

Security Games in Network Flow Problems

Mathieu Dahan

Center for Computational Engineering (CCE)
Massachusetts Institute of Technology (MIT)
Cambridge, MA 02139, mdahan@mit.edu

Saurabh Amin

Department of Civil & Environmental Engineering
Massachusetts Institute of Technology (MIT)
Cambridge, MA 02139, amins@mit.edu

This article considers a two-player strategic game for network routing under link disruptions. Player 1 (defender) routes flow through a network to maximize her value of effective flow while facing transportation costs. Player 2 (attacker) simultaneously disrupts one or more links to maximize her value of lost flow but also faces cost of disrupting links. Linear programming duality in zero-sum games and the Max-Flow Min-Cut Theorem are applied to obtain properties that are satisfied in any Nash equilibrium. A characterization of the support of the equilibrium strategies is provided using graph-theoretic arguments. Finally, conditions under which these results extend to budget-constrained environments are also studied. These results extend the classical minimum cost maximum flow problem and the minimum cut problem to a class of security games on flow networks.

Key words: network security game, flow network, linear programming duality, budget-constrained game.

MSC2000 subject classification: Primary: 91A05, 91A10, 91A43; secondary: 90C46, 05C21

OR/MS subject classification: Primary: games/group decisions: noncooperative; secondary: networks/graphs, programming: linear

1. Introduction. This article studies a class of network security games on flow networks in which simultaneous link disruptions are caused by a strategic adversary. Our setup is motivated by applications in transportation networks where the operator (defender) is interested in strategically routing network flows in the face of adversarial disruptions. We introduce a realistic model of an attacker-defender game played on a network and investigate the physical properties of the game in equilibrium. We also relate the structure of player strategies to solutions of the classical minimum cost maximum flow problem and the minimum cut problem.

Specifically, we study a two-player non-cooperative game over a directed network in which Player 1 (defender or operator) chooses a flow to be routed from a source node to a destination node, and Player 2 (attacker or interdictor) chooses to disrupt one or more edges. In our model, we consider the value of the effective flow and the transportation cost for the defender, and the value of the lost flow and the cost of attack for the attacker. Player 1's payoff linearly increases in the amount of effective flow that reaches the destination node, but decreases with the cost of transporting the initial flow chosen by her. Player 2's payoff linearly increases in the amount of lost flow as a result of an attack chosen by her, and decreases with the cost of conducting the attack. This two-player game is not a zero-sum game. The payoff structures are motivated by the previous formulations in both network interdiction problems (see Wood and Kevin [26], Cormican et al. [10], Bertsimas et al. [9], Avenhaus and Canty [4], Neumayer et al. [21], Ball et al. [6], Ratliff et al. [22], Wollmer [25], Sullivan and Cole Smith [23]), and network security games (see Gueye et al. [16], Szeto [24], Baykal-Grsoy et al. [7]).

Network interdiction problems have already been widely studied in the literature. Our focus is to extend these formulations to simultaneous game settings. Related to our approach is the article by Washburn and Wood [3] in which the authors consider a sequential game, where the operator (leader) chooses one $s - t$ path and then the interdicator (follower) inspects one arc. In this sequential game model, the objective is to maximize the probability with which the operator is detected by the interdicator. In contrast, our simultaneous game formulation captures each player’s strategic uncertainty about her opponent. We also account for the attacker’s cost of attack as well as the defender’s cost of transporting flow through the network. Secondly, we allow both players to have a much larger set of actions. The defender’s feasible flow may contain multiple $s - t$ paths and loops, and the attacker can simultaneously disrupt several edges.

The work by Avenhaus and Canty [4] presents several models of inspection games. One of them considers a two-player game between a passenger of a subway system and the local transit authority. The passenger chooses whether she pays the ticket or not, and simultaneously the transit authority decides whether to inspect or not. This model was motivated by the practicality that inspecting all the time the passenger is either too costly or not worthwhile. Actually, this model is analogous to our game played on a single link (where the transit authority is an interdicator). By adjusting the parameters, we find that these models are strategically equivalent, and our results applied to a single link coincide with the aforementioned inspection game. Thus, our model can be viewed as a generalization of this inspection game to a network setting, where the interdicator chooses the probability of inspecting specific locations or links on the network.¹

Another related line of work in network interdiction games is by Bertsimas et al. [8]. In this sequential game, the operator first chooses a feasible flow, and then the interdicator disrupts a fixed number of edges. The interdicator’s goal is to minimize the largest amount of flow that reaches the destination node. The authors consider two different models of disruption: an arc-based formulation where the flow can be rerouted after an attack, and a path-based formulation where the flow carried by a disrupted edge is lost (or cannot be rerouted). Our formulation is related to the path-based formulation. Since we model a simultaneous game, it is reasonable to assume that the flow through disrupted edges is lost and cannot be re-routed. Although, in Bertsimas et al. [8], the interdicator can disrupt several edges at the same time, she must always disrupt the same number of edges for every action. In our work, we do not assume this restriction on the set of actions for the interdicator.

Our model is also related to the work of Hong and Wooders [17] and Gueye et al. [16] (also see [19]). In these papers, the authors model simultaneous attacker-defender games, where the defender (operator) chooses a feasible flow and the attacker (interdicator) disrupts edges of the network, preventing the flow from reaching the destination node. The major differences with Gueye et al. [16] is that their attacker can only disrupt one edge, and they consider an uncapacitated graph with given supplies and demands, while we consider a capacitated graph with no constraint on the supplies and demands.

We note that Goyal and Vigier [15], and Acemoglu et al. [1] also studied network security from a game-theoretic perspective. Indeed, in their models, the attacker targets the nodes of a network previously chosen by the defender. The defender chooses the network and the allocation of defense resources and the attacker subsequently targets the nodes. The defender’s objective is to minimize the impact of the attack, including network cascades. Since our model assumes that the flow that was supposed to take an edge that is disrupted by an attack is simply lost, it does not consider contagion or dynamic failure propagation through the network. Still, our results illustrate how physical properties such as the expected transportation cost and the expected value of effective flow vary with the attacker’s valuation of lost flow or cost of attack.

¹ However, our set-up does not directly apply to the broader class of inspection games considered by Avenhaus, Von Stengel and Zamir [5], as these games account for strategic manipulation of data by an inspectee who wants to evade detection by a resource constrained inspector.

For the sake of simplicity, we restrict the class of networks we study in this article (even though our ideas apply to a much larger class of networks). This enables us to develop a rather complete characterization of the equilibria of our game. Our first result is that, given the characteristics such as transportation cost and cost of attack, one can tractably compute a Nash equilibrium that is based on minimum cost maximum flows for Player 1 and on minimum cut sets for Player 2. This result is along the lines of previous results by Avenhaus and Canty [4] and of Hong and Wooders [17]. This solution illustrates the key factors that affect the operator’s equilibrium routing strategies in the face of adversarial link disruptions.

We also present theoretical properties satisfied by all the mixed Nash equilibria of our game. A most interesting property is that each player has a unique payoff value in *all* equilibria. We derive analytical expressions of the values of effective and lost flow, and the costs of transportation and attack, in terms of the parameters of the game. These characteristics, common for all equilibria, are derived using a combination of game-theoretic and network optimization ideas, one of them being the Max-Flow Min-Cut Theorem², and do not need full enumeration of the equilibria.

We show that graph-theoretic properties of the network can be used to restrict the support of the strategies that can potentially be Nash equilibria. Recall that computing Nash equilibria of a bimatrix problem is hard. While Daskalakis et al. [11] showed that the computation of a Nash equilibrium is PPAD-complete for a general two-player game, Von Neumann [20] showed that, for a zero-sum two-player game, finding Nash equilibria is equivalent to solving a linear programming problem. Our game is strategically equivalent to a zero-sum game, thus we can use linear programming techniques to compute a Nash equilibrium efficiently. Furthermore, we give an alternative to computing a Nash equilibrium based on a minimum cost maximum flow for the defender and on a minimum cut set for the attacker. Thus, this equilibrium can be computed efficiently by viewing the minimum cost maximum flow problem as a minimum cost circulation problem (see Goldberg and Tarjan [14]). This shows how we can use and extend classical network routing problems to adversarial environments.

Lastly, we study a generalization of the game where both players face budget constraints. Specifically, we view the transportation cost of a flow (resp. cost of an attack) as a resource that needs to be available to the defender (resp. attacker) to send that flow (resp. lead that attack). We compute the minimum budget that players must have to ensure that the equilibrium properties derived for the previous game still hold for the new game defined on action sets with budget constraints. Using the infiniteness of the defender’s set of actions, we give a tight lower bound for the defender’s budget for transporting flows. However, the attacker’s set of actions is discrete and we cannot derive an analogous lower bound for the budget needed by the attacker to disrupt edges. We focus our attention to mixed equilibrium strategies supported over partitions of the minimum cut sets. Then we find the equilibrium strategies in this subset that require the lowest cost of attack. It turns out that we can formulate this problem as an integer programming problem whose optimal solution gives a bound (maybe not the best one) on the attacker’s budget for which our analysis holds.

The rest of the paper is organized as follows: In Section 2 we discuss the main assumptions and present our game model. Section 3 presents our main results on the characterization of Nash equilibria of the game and their relations with classical routing problems. Then in Section 4 we extend our results to a budget-constrained game. Lastly, the implications of relaxing some of the modeling assumptions are considered in Section 5, along with a brief discussion on how to extend our results in these settings.

2. Problem formulation.

² It is not the first time that such classical network optimization ideas are utilized in a game-theoretic context. Indeed, Kalai and Zemel [18] used the Max-Flow Min-Cut Theorem to show that the class of totally balanced games is isomorphic to a particular class of games called *flow games*.

2.1. Preliminaries. Consider a capacitated directed graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ where \mathcal{V} (resp. \mathcal{E}) represents the set of nodes (resp. the set of edges) of \mathcal{G} . For each edge $(i, j) \in \mathcal{E}$, let $c_{ij} \in \mathbb{R}^+$ denote its capacity. Let $s \in \mathcal{V}$ denote the source node and $t \in \mathcal{V}$ the destination node. A flow, defined by the function $x : \mathcal{E} \rightarrow \mathbb{R}^+$, can only enter the network from s and leave from t . There is no demand or supply at other nodes. A flow x is said to be feasible if it satisfies flow conservation at each node and if the flow through each edge does not exceed its capacity:

$$\begin{aligned} \forall i \in \mathcal{V} \setminus \{s, t\}, \quad \sum_{(j,i) \in \mathcal{E}} x(j,i) &= \sum_{(i,j) \in \mathcal{E}} x(i,j) \\ \forall (i,j) \in \mathcal{E}, \quad 0 \leq x(i,j) &\leq c_{ij}. \end{aligned}$$

Let \mathcal{F} denote the set of feasible flows, and Λ the set containing all the loops and source-destination $s - t$ paths of the network. Let $x_{ij} := x(i, j)$ denote the flow through the edge (i, j) , and x_λ the quantity of flow of x sent through $\lambda \in \Lambda$. The edge flows x_{ij} and loop/path flows x_λ satisfy:

$$\forall (i, j) \in \mathcal{E}, \quad x_{ij} = \sum_{\{\lambda \in \Lambda \mid (i,j) \in \lambda\}} x_\lambda. \quad (1)$$

An $s - t$ cut is a partition $\{\mathcal{S}, \mathcal{T}\}$ of \mathcal{V} , such that $s \in \mathcal{S}$ and $t \in \mathcal{T}$. The *cut-set* of $\{\mathcal{S}, \mathcal{T}\}$ and its *capacity* are defined as $E(\{\mathcal{S}, \mathcal{T}\}) = \{(i, j) \in \mathcal{E} \mid i \in \mathcal{S}, j \in \mathcal{T}\}$ and $C(\{\mathcal{S}, \mathcal{T}\}) = \sum_{(i,j) \in E(\{\mathcal{S}, \mathcal{T}\})} c_{ij}$. Let $F(x) = \sum_{\{i \in \mathcal{V} \mid (i,t) \in \mathcal{E}\}} x_{it}$ denote the amount of flow passing from the source s to the sink t . We recall the *maximum flow problem*:

$$\begin{aligned} (\mathcal{P}_1) \quad & \text{maximize } F(x) \\ & \text{subject to } x \in \mathcal{F}. \end{aligned}$$

The well-known *Max-Flow Min-Cut Theorem* by Ford and Fulkerson [13] states that the optimal value of the maximum flow problem is equal to the minimum capacity over all $s - t$ cuts. We call *min-cut set* a cut-set with minimum capacity. We also state the *minimum cost maximum flow problem* by Edmonds and Karp [12]:

$$\begin{aligned} (\mathcal{P}_2) \quad & \text{minimize } \sum_{(i,j) \in \mathcal{E}} b_{ij} x_{ij} \\ & \text{subject to } x \in \mathcal{F} \\ & F(x) \geq F(x'), \quad \forall x' \in \mathcal{F}, \end{aligned}$$

where for every edge $(i, j) \in \mathcal{E}$, $b_{ij} \in \mathbb{R}^+$ denotes the cost of transporting a unit flow through (i, j) .

We use F^{\max} (resp. Ω_{\max}) to denote the optimal value (resp. optimal solution set) of the max-flow problem (\mathcal{P}_1) . Similarly, we denote the optimal value (resp. the set of optimal solutions) of problem (\mathcal{P}_2) by T^{\min} (resp. Ω^*). Note that $\Omega^* \subseteq \Omega_{\max}$ (i.e., any min-cost max-flow is a max-flow).

2.2. Model. We define a simultaneous, semi-infinite, two-player strategic game $\Gamma := \langle \{1, 2\}, (\mathcal{F}, \mathcal{A}), (u_1, u_2) \rangle$. Player 1 (**P1**) is the defender (operator) who chooses to route a flow $x \in \mathcal{F}$ through the network, and player 2 (**P2**) is the attacker (interdictor) who chooses an attack μ to disrupt a subset of edges of \mathcal{G} . The action set for **P1** (resp. **P2**) is given by \mathcal{F} (resp. $\mathcal{A} := \{0, 1\}^{\mathcal{E}}$).

An attack μ is a function from \mathcal{E} to $\{0, 1\}$ defined as follows:

$$\mu_{ij} := \mu(i, j) = \begin{cases} 1 & \text{if } (i, j) \text{ is disrupted,} \\ 0 & \text{otherwise.} \end{cases}$$

Note that **P2**'s choice of an attack μ can lead to the disruption of multiple edges of the network.

We use the following notation to describe certain specific player actions: x^0 the action of not sending flow in the network, $x^* \in \Omega^*$ an optimal solution of (\mathcal{P}_2) i.e., a min-cost max-flow, μ^0 the action of not attacking any edge of the network, and μ^{min} the action that disrupts all the edges of a min-cut set of the network.

Since Γ is a simultaneous game, we assume that after an edge is disrupted, the flow that was supposed to cross this edge (if there were no attack) is lost and it is not re-routed. For the sake of simplicity, we do not consider attacks that can only result in partially disrupted edges and might still permit some flow to pass through the attacked edges. Thus, the *effective flow*, denoted x^μ , when a flow x is chosen by **P1** and an attack μ is chosen by **P2** can be expressed as follows:

$$\forall (i, j) \in \mathcal{E}, \quad x_{ij}^\mu = \sum_{\lambda \in \Lambda_{ij}^\mu} x_\lambda,$$

where $\Lambda_{ij}^\mu := \{\lambda \in \Lambda \mid (i, j) \in \lambda \text{ and } \forall (i', j') \in \lambda, \mu_{i'j'} = 0\}$. That is, the effective flow through an edge (i, j) is the sum of all the initial path flows through edge (i, j) that do not contain any attacked edge. The effective flow x^μ can be viewed as a feasible flow in \mathcal{F} that successfully carries the amount of flow from x that is not lost due to the attack μ .

In this model, the payoff of **P1** is defined as the value of effective flow assessed by **P1** net the cost of transporting the initial flow:

$$u_1(x, \mu) = \underbrace{p_1 F(x^\mu)}_{\text{value of effective flow}} - \underbrace{T(x)}_{\text{transportation cost}} \quad (2)$$

where $p_1 \in \mathbb{R}^+$ is the marginal value of the flow for **P1**, and $T(x) := \sum_{(i,j) \in \mathcal{E}} b_{ij} x_{ij}$ is the cost of transporting the initial flow x . Thus, when one additional unit of flow reaches t , **P1**'s payoff increases by p_1 and at the same time decreases by its transportation cost.

Similarly, the payoff of **P2** is defined as the value of lost flow assessed by **P2** net the cost of executing the attack:

$$u_2(x, \mu) = \underbrace{p_2 F(x - x^\mu)}_{\text{value of lost flow}} - \underbrace{C(\mu)}_{\text{cost of attack}} \quad (3)$$

where $p_2 \in \mathbb{R}^+$ is the marginal value of the lost flow for **P2** (in general, $p_1 \neq p_2$), and $C(\mu) := \sum_{(i,j) \in \mathcal{E}} c_{ij} \mu_{ij}$ is the cost of the attack μ . Thus, if the disruption of an edge induces the loss of one unit of flow, the payoff of **P2** increases by p_2 , and at the same time decreases by the cost of attack.

In this model, we suppose that the cost of attacking an edge is proportional to its capacity. Note that after rescaling **P2**'s payoff, without loss of generality, the cost of attacking an edge becomes equal to its capacity. Also, notice that F and T are linear forms on \mathcal{F} , and C is a linear form on \mathcal{A} . However, by definition, \mathcal{F} and \mathcal{A} are not vector spaces.

We allow each player to randomize over her set of pure actions. Let $\Delta(\mathcal{F})$ and $\Delta(\mathcal{A})$ denote the mixed extensions of **P1**'s and **P2**'s pure strategies, respectively, i.e.:

$$\Delta(\mathcal{F}) = \left\{ \sigma^1 \in [0, 1]^{\mathcal{F}} \mid \sum_{x \in \mathcal{F}} \sigma^1(x) = 1 \right\}, \quad \Delta(\mathcal{A}) = \left\{ \sigma^2 \in [0, 1]^{\mathcal{A}} \mid \sum_{\mu \in \mathcal{A}} \sigma^2(\mu) = 1 \right\}.$$

For notational simplicity, we define $\sigma_x^1 := \sigma^1(x)$ and $\sigma_\mu^2 := \sigma^2(\mu)$. Given any function $\varphi : \mathcal{F} \times \mathcal{A} \rightarrow \mathbb{R}$ and a mixed strategy profile $\sigma = (\sigma^1, \sigma^2) \in \Delta(\mathcal{F}) \times \Delta(\mathcal{A})$, we denote $\mathbb{E}_\sigma[\varphi(x, \mu)] := \sum_{x \in \mathcal{F}} \sigma_x^1 \sum_{\mu \in \mathcal{A}} \sigma_\mu^2 \varphi(x, \mu)$ the expectation of φ with respect to σ .

Given a strategy profile $\sigma = (\sigma^1, \sigma^2) \in \Delta(\mathcal{F}) \times \Delta(\mathcal{A})$, the respective player expected payoffs can be expressed as:

$$U_1(\sigma^1, \sigma^2) = p_1 \mathbb{E}_\sigma [F(x^\mu)] - \mathbb{E}_\sigma [T(x)] \quad (4)$$

$$U_2(\sigma^1, \sigma^2) = p_2 (\mathbb{E}_\sigma [F(x)] - \mathbb{E}_\sigma [F(x^\mu)]) - \mathbb{E}_\sigma [C(\mu)]. \quad (5)$$

We will use the notation $U_i(x, \sigma^2) = U_i(\mathbb{1}_{\{x\}}, \sigma^2)$ and $U_i(\sigma^1, \mu) = U_i(\sigma^1, \mathbb{1}_{\{\mu\}})$ for $i \in \{1, 2\}$. Thus, the mixed extension of the game Γ is $\langle \{1, 2\}, (\Delta(\mathcal{F}), \Delta(\mathcal{A})), (U_1, U_2) \rangle$.

Let us illustrate this model through an example.

EXAMPLE 1. Consider the network shown in Fig. 1a. The edge labels give the capacities and transportation costs. Both players play one shot of the game according to Fig. 1b.

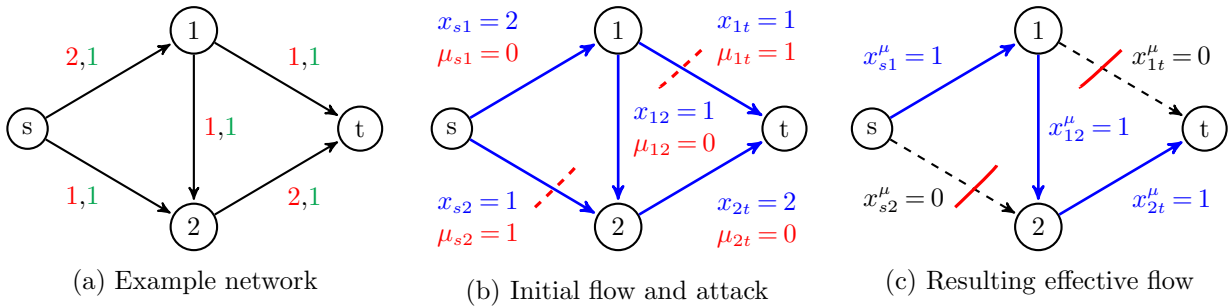


FIGURE 1. The labels of each edge in Fig. 1a correspond to its capacity (red) and its transportation cost (green).

In this example, **P1** sends one unit of flow through each of the paths $\{s, 1, t\}$, $\{s, 1, 2, t\}$ and $\{s, 2, t\}$, and **P2** disrupts edges $(1, t)$ and $(s, 2)$. Therefore, the flows through paths $\{s, 1, t\}$ and $\{s, 2, t\}$ are lost and the effective flow, shown in Fig. 1c, consists of the unit flow through the path $\{s, 1, 2, t\}$, i.e., $F(x^\mu) = 1$. Since that each edge (i, j) has a transportation cost $b_{ij} = 1$, the cost of transporting the initial flow x is $T(x) = 7$. Thus, **P1**'s payoff is $u_1(x, \mu) = p_1 - 7$.

The amount of lost flow $F(x - x^\mu)$ that results from the attack μ is equal to 2. Since **P2** disrupted 2 edges of capacity 1 each, the cost of attack $C(\mu) = 2$. Thus, **P2**'s payoff is $u_2(x, \mu) = 2p_2 - 2$.

2.3. Standard definitions and main assumption. We recall the following standard definitions:

The *support* of $\sigma^1 \in \Delta(\mathcal{F})$ (resp. $\sigma^2 \in \Delta(\mathcal{A})$) is $\text{supp}(\sigma^1) = \{x \in \mathcal{F} \mid \sigma_x^1 > 0\}$ (resp. $\text{supp}(\sigma^2) = \{\mu \in \mathcal{A} \mid \sigma_\mu^2 > 0\}$).

A mixed strategy profile $(\sigma^{1*}, \sigma^{2*}) \in \Delta(\mathcal{F}) \times \Delta(\mathcal{A})$ is a Nash Equilibrium (NE) if and only if:

$$\forall \sigma^1 \in \Delta(\mathcal{F}), U_1(\sigma^{1*}, \sigma^{2*}) \geq U_1(\sigma^1, \sigma^{2*}), \quad (6)$$

$$\forall \sigma^2 \in \Delta(\mathcal{A}), U_2(\sigma^{1*}, \sigma^{2*}) \geq U_2(\sigma^{1*}, \sigma^2). \quad (7)$$

We denote Σ the set of NE of the game Γ . Equivalently, at a NE $(\sigma^{1*}, \sigma^{2*})$, σ^{1*} (resp. σ^{2*}) is a Best Response (BR) to σ^{2*} (resp. σ^{1*}).

A two-player game is a *strictly competitive game* (SCG) if, when both players change their mixed strategies, either the expected payoffs remain the same or one of the expected payoffs strictly increases and the other strictly decreases. In particular, a *zero-sum* game (i.e., $u_1 = -u_2$) is an SCG. Adler et al. [2] define SCG using the notion of affine variance: u_1 is an affine variant of $-u_2$ if and only if $\exists (\lambda, \beta) \in \mathbb{R}_+^* \times \mathbb{R} \mid \forall (x, \mu) \in \mathcal{F} \times \mathcal{A}, u_1(x, \mu) = -\lambda u_2(x, \mu) + \beta$. The game Γ is an SCG if and only if u_1 is an affine variant of $-u_2$.

Consider two games $\Gamma = \langle \{1, 2\}, (\mathcal{F}, \mathcal{A}), (u_1, u_2) \rangle$, $\tilde{\Gamma} = \langle \{1, 2\}, (\mathcal{F}, \mathcal{A}), (\tilde{u}_1, \tilde{u}_2) \rangle$ with:

$$\begin{aligned}\tilde{u}_1(x, \mu) &= a_1 u_1(x, \mu) + g(\mu) \\ \tilde{u}_2(x, \mu) &= a_2 u_2(x, \mu) + h(x)\end{aligned}$$

where $(a_1, a_2) \in (\mathbb{R}_+^*)^2$, $g: \mathcal{A} \rightarrow \mathbb{R}$ and $h: \mathcal{F} \rightarrow \mathbb{R}$. Then, Γ and $\tilde{\Gamma}$ are strategically equivalent (i.e., they have the same set of equilibria).

Our objective in this article is to develop a complete characterization of the equilibria of the game Γ and to relate them to the solutions of classical routing problems. To facilitate this, we restrict our attention to the class of networks that satisfy the following assumption:

ASSUMPTION 1. *Let $\alpha := \min_{\lambda \in \Lambda_{path}} \sum_{(i,j) \in \lambda} b_{ij}$. There exists an optimal solution of (\mathcal{P}_2) , $x^* \in \Omega^*$, that takes $s-t$ paths with marginal transportation cost equal to α , i.e.,*

$$\exists x^* \in \Omega^* \text{ s.t. } \forall \lambda \in \Lambda_{path} : x_\lambda > 0 \implies \sum_{(i,j) \in \lambda} b_{ij} = \alpha,$$

where Λ_{path} is the set containing all the $s-t$ paths of the network.

This assumption, noted (A1), implies that if $x^* \in \Omega^*$ denotes a min-cost max-flow, the cost of transporting a unit flow through each $s-t$ path taken by x^* is identically equal to α . By definition of α , every other path in the network cannot have a smaller marginal transportation cost. Notice that if such an x^* exists, then (A1) will be satisfied for any optimal solution of (\mathcal{P}_2) . The case when every $s-t$ has an identical marginal transportation cost is a special case of this assumption. In Section 5.1, we will discuss the implications of relaxing (A1). We illustrate (A1) with an example.

EXAMPLE 2. Consider the network flow problem in Fig. 2. There is a unique min-cost max-flow x^* , which carries 1 unit of flow through paths $\{s, 2, 4, t\}$, $\{s, 2, 3, t\}$ and $\{s, 1, t\}$. Thus, the total amount of flow is equal to 3 units. In this network, $\alpha = 3$, and each path taken by x^* has a marginal transportation cost equal to 3. Thus, the cost of transporting x^* is equal to 9. The remaining paths that are not taken by x^* are $\{s, 4, t\}$ with a transportation cost 4, and $\{s, 1, 3, t\}$ with a transportation cost 3. Thus, (A1) is satisfied.




FIGURE 2. Min-cost max-flow (drawn in bold blue) in a network satisfying (A1). The labels of each edge correspond to the flow it carries (blue), its capacity (red) and its transportation cost (green).

REMARK 1. (A1) implies that for all $x \in \mathcal{F}$, $T(x) \geq \alpha F(x)$ and $T^{\min} = \alpha F^{\max}$. To show this, we note b_λ the cost of transporting one unit of flow through path $\lambda \in \Lambda_{path}$, then:

$$\forall x \in \mathcal{F}, T(x) = \sum_{\lambda \in \Lambda} b_\lambda x_\lambda \geq \sum_{\lambda \in \Lambda_{path}} b_\lambda x_\lambda \stackrel{(A1)}{\geq} \alpha \sum_{\lambda \in \Lambda_{path}} x_\lambda = \alpha F(x), \quad (8)$$

and

$$T^{\min} = T(x^*) = \sum_{\lambda \in \Lambda_{path}} b_\lambda x_\lambda^* \stackrel{(A1)}{=} \alpha \sum_{\lambda \in \Lambda_{path}} x_\lambda^* = \alpha F(x^*) = \alpha F^{\max}, \quad (9)$$

where we used the fact that any min-cost max-flow does not send flow in any loop.

Note that our setup can be easily extended to networks with multiple sources and multiple destination nodes, but still satisfying (A1). For such a network, one needs to add an extra source (resp. destination) node and connect it to every existing source (resp. destination) node with an uncapacitated edge of cost of transportation equal to zero. This modification gives a new network with single source and single destination. The outcome of the game defined for the original network remains the same as that of the game defined for the new network.

Next, we focus on the NE of the game Γ , and utilize the properties of NE to relate the support of equilibrium strategies of **P1** (resp. **P2**) with the solutions of (\mathcal{P}_2) (resp. min-cut sets).

3. Properties of Nash equilibria. In this section, we present theoretical and practical properties satisfied by the NE of the game Γ . First, we give one NE of Γ that is based on min-cost max-flows for **P1** and on min-cut sets for **P2**. Next, we focus on the parameter range of interest $p_1 > \alpha$, $p_2 > 1$, and derive analytical expressions of certain physical quantities of interest at any NE. This involves a combination of game-theoretic arguments and network optimization results. Finally, we show how we can restrict the support of the strategies that can be NE using graph-theoretic properties of the network.

3.1. Preliminary results. The following lemma states that although Γ is not a zero-sum game, it is strategically equivalent to a zero-sum game.

LEMMA 1. Γ is strategically equivalent to $\tilde{\Gamma} := \langle \{1, 2\}, (\mathcal{F}, \mathcal{A}), (\tilde{u}_1, -\tilde{u}_1) \rangle$ where:

$$\forall (x, \mu) \in \mathcal{F} \times \mathcal{A}, \quad \tilde{u}_1(x, \mu) = F(x^\mu) - \frac{1}{p_1} T(x) + \frac{1}{p_2} C(\mu). \quad (10)$$

Therefore, the NE of Γ can be obtained by solving the following two linear programming problems:

$$\begin{array}{l|l} (LP_1) \text{ maximize} & z \\ \text{subject to} & \tilde{U}_1(\sigma^1, \mu) \geq z, \forall \mu \in \mathcal{A} \\ & \sigma^1 \in \Delta(\mathcal{F}) \end{array} \quad \left| \quad \begin{array}{l} (LP_2) \text{ maximize} & z' \\ \text{subject to} & \tilde{U}_2(x, \sigma^2) \geq z', \forall x \in \mathcal{F} \\ & \sigma^2 \in \Delta(\mathcal{A}) \end{array} \right.$$

If $(\sigma^{1*}, \sigma^{2*}) \in \Sigma$ and $(\sigma^{1\dagger}, \sigma^{2\dagger}) \in \Sigma$, then $(\sigma^{1*}, \sigma^{2\dagger}) \in \Sigma$ and $(\sigma^{1\dagger}, \sigma^{2*}) \in \Sigma$ (interchangeability). Furthermore, Σ is a convex set.

Proof in Appendix A.1.

The following lemma states that **P1**'s all strategies containing loops are strictly dominated.


LEMMA 2. Any flow containing loops is not a BR for **P1**.

Proof in Appendix A.1.

The intuition behind this result is that if **P1** sends flow in a loop, then she will pay an extra cost without increasing the amount of flow that can reach the terminal node. Therefore, **P1** has no incentive to send flow in any loop. Thus, Λ in (1) can be restricted to the set of $s - t$ paths, and \mathcal{F} can be restricted to the set of feasible flows that do not take any loop.

Props. 1–3 below provide that, for given p_1 , p_2 and α (given and fixed under (A1)), the game Γ admits qualitatively different equilibria in regions $0 < p_1 < \alpha$ and $p_2 > 0$ (Region **I**), $p_1 > \alpha$ and $0 < p_2 < 1$ (Region **II**), and $p_1 > \alpha$ and $p_2 > 1$ (Region **III**). The proofs are provided in Appendix A.1. These regions are illustrated in Fig. 3.

The following result states that no flow and no attack is the unique NE of Γ in Region **I**:

FIGURE 3. Support of equilibrium strategies in Regions **I-III**.

PROPOSITION 1 (Region I). *If $p_1 < \alpha$, then $\Sigma = \{(x^0, \mu^0)\}$, with $u_1(x^0, \mu^0) = 0$ and $u_2(x^0, \mu^0) = 0$.*

Intuitively, when $0 < p_1 < \alpha$, the marginal value of effective flow that reaches the destination node t is less than the marginal transportation cost for every $s-t$ path. Therefore, **P1** will face negative utility if she sends flow through the network. Thus, in this case, her BR is not to route any flow. Since no flow is sent by **P1**, **P2**'s BR is not to attack, otherwise she will face the cost of attack without gaining any value from lost flow.

Next, for Region **II**, we obtain that min-cost max-flow and no attack is a pure NE.

PROPOSITION 2 (Region II). *If $p_1 > \alpha$ and $p_2 < 1$, then $\forall x^* \in \Omega^*$, $\{x^*, \mu^0\} \in \Sigma$. The equilibrium payoffs are $u_1(x^*, \mu^0) = (p_1 - \alpha)F^{max}$ and $u_2(x^*, \mu^0) = 0$.*

This result can be explained as follows: on one hand, since **P2**'s valuation of lost flow is small ($p_2 < 1$), for any attack, the utility gained from the lost flow is always lower than the cost of attack. Therefore, **P2**'s BR is not to attack any edge. On the other hand, **P1**'s valuation of effective flow reaching t is higher than the disutility it faces in transportation costs ($p_1 > \alpha$). Since **P2** does not disrupt any edge, every flow sent through the network reaches t ; thus, **P1**'s BR is to send a maximum flow. Among the different maximum flows, a min-cost max-flow maximizes **P1**'s equilibrium payoff. Note that if $p_1 = \alpha$ and $p_2 < 1$, then both (x^0, μ^0) and (x^*, μ^0) are NE. The equilibrium payoffs are still $(0, 0)$.

The following Proposition 3 shows that in Region **III**, Γ admits a NE whose support is based on a min-cost max-flow for **P1**, and on a min-cut set for **P2**.

PROPOSITION 3 (Region III). *If $p_1 > \alpha$ and $p_2 > 1$, then Γ has no pure NE. Furthermore, $\exists \tilde{\sigma} = (\tilde{\sigma}^1, \tilde{\sigma}^2) \in \Sigma$ such that $U_1(\tilde{\sigma}^1, \tilde{\sigma}^2) = U_2(\tilde{\sigma}^1, \tilde{\sigma}^2) = 0$, and $\text{supp}(\tilde{\sigma}^1) = \{x^0, x^*\}$ and $\text{supp}(\tilde{\sigma}^2) = \{\mu^0, \mu^{min}\}$. The corresponding probabilities are given by:*

$$\tilde{\sigma}_{x^0}^1 = 1 - \frac{1}{p_2}, \quad \tilde{\sigma}_{x^*}^1 = \frac{1}{p_2}, \quad (11)$$

$$\tilde{\sigma}_{\mu^0}^2 = \frac{\alpha}{p_1}, \quad \tilde{\sigma}_{\mu^{min}}^2 = 1 - \frac{\alpha}{p_1}. \quad (12)$$

We can now make a few useful observations. First, in contrast to Props. 1 and 2, in Region **III**, both players necessarily randomize their actions in any equilibrium. Indeed, if **P1**'s pure strategy is to route some flow x in the network, then **P2**'s BR is to disrupt some edges taken by x in order to induce the maximum loss. Then **P1** has an incentive to change her strategy and route another flow that takes other paths not disrupted by **P2**. Therefore, for every pure action profile, at least one of the players has an incentive to deviate, preventing them from reaching a pure strategy equilibrium.

Second, the mixed equilibrium $(\tilde{\sigma}^1, \tilde{\sigma}^2)$, as defined in (11) and (12), can be obtained from a solution x^* of problem (\mathcal{P}_2) and a min-cut set of the graph \mathcal{G} . Note that $(\tilde{\sigma}^1, \tilde{\sigma}^2)$ has a simple structure: **P1** either sends a min-cost max-flow, or does not send any flow in the network. Similarly, **P2** either disrupts all the edges of a min-cut set, or does not attack any edge of the network.

Finally, Prop. 3 provides a game-theoretic intuition: **P1**'s equilibrium strategy $\tilde{\sigma}^1$ is characterized by p_2 , and similarly, **P2**'s equilibrium strategy $\tilde{\sigma}^2$ is characterized by p_1 and α . This can be explained as follows: as p_2 increases, $\tilde{\sigma}_{x^*}^1$ decreases while $\tilde{\sigma}_{x_0}^1$ increases. When **P2**'s valuation of lost flow is large, she has more incentive to attack, so any flow sent by **P1** will be more likely to be lost. Thus, **P1** chooses not to send any flow with higher probability than sending x^* . Likewise, as p_1 increases, $\tilde{\sigma}_{\mu^{min}}^2$ increases while $\tilde{\sigma}_{\mu^0}^2$ decreases. Again, when the marginal valuation of effective flow is large, **P1** will prefer to send as much flow as she can. Thus, **P2** will be more likely to attack a min-cut set.

The following example applies the results of Props. 1-3:

EXAMPLE 3. Consider the network in Fig. 4. We can see that $\alpha = 3$, and that the min-cost max-flow sends 1 unit of flow through $\{s, 1, 3, t\}$, $\{s, 2, 3, t\}$ and $\{s, 2, 4, t\}$, and only takes paths with transportation cost equal to 3. Thus, (A1) is satisfied. The min-cut set is given by $\{(1, 3), (2, 3), (2, 4)\}$. The NE described in Props. 1-3 are illustrated in Fig. 5.




FIGURE 4. Example network. Edge capacities and transportation costs are labeled in red and green colors respectively.

Thanks to Props. 1-3, each player has a strategy that achieves a stable outcome for any p_1 and p_2 , and any network satisfying (A1). However, **P1** may be interested in an equilibrium strategy that maximizes the expected amount of effective flow that successfully crosses the network after attack. Similarly, **P2** may be interested in an equilibrium strategy that maximizes the expected amount of lost flow. Moreover, players may be subject to certain budget constraints. To answer these questions, we need to further analyze the set of NE. Regions I and II do not require further study; however, Region III hosts many useful properties which we present next.

3.2. Main theorem.

THEOREM 1. If $p_1 > \alpha$, $p_2 > 1$, and under (A1), then for any $\sigma^* \in \Sigma$:

(i) Both players' equilibrium payoffs are equal to 0, i.e.:

$$U_1(\sigma^{1*}, \sigma^{2*}) \equiv 0 \quad (13)$$

$$U_2(\sigma^{1*}, \sigma^{2*}) \equiv 0 \quad (14)$$

(ii) The expected amount of flow sent in the network is given by:

$$\mathbb{E}_{\sigma^*} [F(x)] \equiv \frac{1}{p_2} F^{max} \quad (15)$$

and the expected transportation cost is given by:

$$\mathbb{E}_{\sigma^*} [T(x)] \equiv \frac{1}{p_2} T^{min} \quad (16)$$




FIGURE 5. NE described in Props. 1, 2 and 3. The min-cost max-flow (resp. min-cut set attack) is in bold blue (resp. dotted red).

(iii) The expected cost of attack is given by:

$$\mathbb{E}_{\sigma^*} [C(\mu)] \equiv F^{\max} - \frac{1}{p_1} T^{\min} \quad (17)$$

(iv) The expected amount of effective flow (that reaches t) is given by:

$$\mathbb{E}_{\sigma^*} [F(x^\mu)] \equiv \frac{1}{p_1 p_2} T^{\min}. \quad (18)$$

Proof in Appendix A.2.

From Thm. 1 we observe that, in any equilibrium for Region **III**, the expected amount of initial and effective flow and expected transportation cost to **P1**, and the expected cost of attack to **P2**, can be computed in closed form using the parameters p_1 , p_2 , F^{\max} , and T^{\min} . It is easy to check that Thm. 1 is satisfied by $(\tilde{\sigma}^1, \tilde{\sigma}^2)$ defined in (11) and (12).

Interestingly, the payoffs of both players are zero for any NE (ref. (13) and (14)). Note that in general, a game that is strategically equivalent to a zero-sum game can have different NE that lead to different payoffs. Thus, (i) in Thm. 1 cannot be entirely derived from the equivalent zero-sum game $\tilde{\Gamma}$ and requires the application of other results such as the Max-Flow Min-Cut Theorem (see the proof of Thm. 1). Now we further explain (ii)-(iv) in Thm. 1.

Following (15) (resp. (16)), the expected amount of initial flow (resp. the expected cost of transportation) at equilibrium is equal to some fraction of the value (resp. transporting cost) of the min-cost max-flows, and these expectations decrease with p_2 .

Following (17), the expected cost of attack at any NE is a constant. In fact, under (A1), we know that $T^{\min} = \alpha F^{\max}$ (see (9)). Therefore, the expected cost of attack at equilibrium becomes $(1 - \frac{\alpha}{p_1}) F^{\max}$. Applying the Max-Flow Min-Cut Theorem, we obtain that the expected cost of

attack is equal to some fraction of the cost of attacking a min-cut set, and this fraction increases with p_1 .

Following (18), the expected amount of effective flow is again a constant at equilibrium. Under (A1), this quantity becomes $\frac{\alpha}{p_1 p_2} F^{\max}$. Although the amount of effective flow depends on both players' strategies, in any NE, its expectation is always equal to some fraction of the amount of max-flow. Since $\frac{\alpha}{p_1 p_2} < \frac{1}{p_2}$, this flow is always smaller than the expected amount of initial flow. Interestingly, the expected amount of effective flow *decreases* when p_1 and/or p_2 increase. This result can be explained by noting that when p_1 increases, the disruption caused by **P2** increases, so there is more lost flow and the expected effective flow decreases.

Thm. 1 also enables estimation of the expected amount of lost flow and the *yield* of **P1** in any NE. We define yield as the ratio of the expected amount of effective flow and the expected amount of initial flow in the network. We have the following corollary:

COROLLARY 1. *For any $\sigma^* \in \Sigma$, the expected amount of lost flow is given by:*

$$\mathbb{E}_{\sigma^*} [F(x - x^\mu)] \equiv \frac{1}{p_2} \left(F^{\max} - \frac{1}{p_1} T^{\min} \right) \stackrel{(9)}{=} \frac{1}{p_2} \left(1 - \frac{\alpha}{p_1} \right) F^{\max}, \quad (19)$$

and the yield is given by:

$$\frac{\mathbb{E}_{\sigma^*} [F(x^\mu)]}{\mathbb{E}_{\sigma^*} [F(x)]} \equiv \frac{T^{\min}}{p_1 F^{\max}} \stackrel{(9)}{=} \frac{\alpha}{p_1}. \quad (20)$$

From (19), in any equilibrium, the expected amount of lost flow is equal to some fraction of the amount of max-flow. The corresponding coefficient increases with p_1 , because when p_1 is large, **P1** sends more flow and **P2** disrupts more edges. However, the coefficient decreases when p_2 increases, because when p_2 is large, **P2** causes more disruption and **P1** sends less flow in the network. Finally, from (20), in any equilibrium, the yield decreases in the ratio $\frac{p_1}{\alpha}$, but it *does not* depend on p_2 or the maximum amount of flow F^{\max} . When $\frac{p_1}{\alpha}$ is large, **P1** has more incentive to send flow in the network and **P2** will attack more frequently, resulting in a lower yield of the network.

Thus, Thm. 1 provides many properties that are satisfied by any NE in Region **III**. Next, we study the support of NE and relate it to the optimal solutions of (\mathcal{P}_2) and the min-cut sets.

3.3. Necessary conditions. Recall the NE in Prop. 3 which has a support based on a min-cost max-flow (for **P1**) and on a min-cut set (for **P2**). We now investigate the generality of this result to other NE. This leads to additional properties satisfied by all NE, and in many cases also eases the computation of NE.

First, let us present a result regarding the paths taken by the flows in the support of **P1**'s strategy at equilibrium.

LEMMA 3. *Under (A1), every flow in the support of a NE only takes paths whose marginal transportation cost is α .*

$$\forall (\sigma^{1*}, \sigma^{2*}) \in \Sigma, \forall x \in \text{supp}(\sigma^{1*}), \forall \lambda \in \Lambda, x_\lambda > 0 \implies \sum_{(i,j) \in \lambda} b_{ij} = \alpha.$$

Proof in Appendix A.3.

In other words, the paths that induce a transportation cost strictly greater than α are not chosen in equilibrium. Notice that this lemma implies that any max-flow that is not a min-cost max-flow is not in the support of any NE. This lemma is useful for constructing **P1**'s equilibrium strategies when the network has only a small number of paths of marginal transportation cost α . In the case when the only paths of marginal transportation cost equal to α are the ones taken by a unique

min-cost max-flow x^* , we can deduce that all the equilibrium strategies of **P1** can be constructed from x^* (or sub-flows of x^*). This lemma is of limited use in the construction of **P1**'s equilibrium strategies if most of the paths of the network have the same smallest transportation cost.

Secondly, the following result characterizes the support of **P2**'s equilibrium strategies.

PROPOSITION 4. *Every attack in the support of a NE has a cost at most equal to the cost of attacking a min-cut set, and disrupts edges that are saturated by every min-cost max-flow:*

$$\begin{aligned} \forall(\sigma^{1*}, \sigma^{2*}) \in \Sigma, \forall \mu \in \text{supp}(\sigma^{2*}): C(\mu) \leq C(\mu^{\min}) = F^{\max} \\ \forall(i, j) \in \mathcal{E}, \mu_{ij} = 1 \implies \forall x^* \in \Omega^*, x_{ij}^* = c_{ij}. \end{aligned} \quad (21)$$

Proof in Appendix A.3.

This result tells us that the attacks that require a cost that is strictly greater than the cost of disrupting a min-cut set are not chosen in any equilibrium. Further, if an edge is not saturated by at least one min-cost max-flow, then it is not disrupted in equilibrium. Recall that **P2**'s set of actions is isomorphic to the power set of \mathcal{E} , that has $2^{|\mathcal{E}|}$ elements which can be huge. Therefore, Prop. 4 drastically restricts the set of actions that can be potentially chosen in equilibrium by **P2**.

However, since the edges that are part of a min-cut set are saturated by every (min-cost) max-flow, we cannot restrict the set of edges that can be potentially disrupted in equilibrium beyond the min-cut sets. Nevertheless, one can find NE in which edges that are not part of any min-cut set are disrupted with positive probability. The following example illustrates this point.

EXAMPLE 4. Consider the network in Fig. 6. We can see that $\alpha = 3$, and the min-cost max-




FIGURE 6. Example network in which edges outside a min-cut set are disrupted in equilibrium.

flow x^* sends 1 unit of flow through $\{s, 2, t\}$, $\{s, 3, 2, t\}$ and $\{s, 3, t\}$, and only takes paths with transportation cost equal to 3. Thus, (A1) is satisfied. In this network, there is a unique min-cut set given by $\{(2, t), (3, t)\}$. Let $\mu' = \mathbb{1}_{\{(s, 2), (s, 3)\}}$ the attack that disrupts edges $(s, 2)$ and $(s, 3)$. In the case when $3 < p_1 < 4$ and $p_2 > 1$, one can see that there exists a NE $(\tilde{\sigma}^1, \sigma^{2*})$ where $\tilde{\sigma}^1$ is defined in (11), and σ^{2*} is defined by $\sigma_{\mu^0}^{2*} = \frac{\alpha}{p_1}$ and $\sigma_{\mu'}^{2*} = 1 - \frac{\alpha}{p_1}$. However, $(s, 2)$ and $(s, 3)$ are not part of the min-cut set.

Finally, since there are equilibrium strategies that disrupt at least one min-cut set (see e.g. Prop. 3), it is useful to estimate the amount of lost flow for each edge that is attacked in a min-cut set. Similarly, the probability with which each edge of a min-cut set will be disrupted is also of interest because it can be interpreted as the probability with which the flow routed by **P1** is lost when **P2**'s equilibrium strategy only involves edges belonging to a min-cut set. The following proposition answers these questions.

PROPOSITION 5. *Consider a min-cut set $E(\{\mathcal{S}, \mathcal{T}\})$, then:*

$$\forall(\sigma^{1*}, \sigma^{2*}) \in \Sigma, \forall(i, j) \in E(\{\mathcal{S}, \mathcal{T}\}), \mathbb{E}_{\sigma^*}[x_{ij}] = \frac{c_{ij}}{p_2}. \quad (22)$$

Furthermore, for any NE $\sigma^* \in \Sigma$ whose support only contains attacks that disrupt edges of $E(\{\mathcal{S}, \mathcal{T}\})$, we have:

$$\forall (i, j) \in E(\{\mathcal{S}, \mathcal{T}\}), \mathbb{P}_{\sigma^*}(\{(i, j) \text{ is disrupted}\}) = 1 - \frac{\alpha}{p_1}. \quad (23)$$

Proof in Appendix A.3.

From (22), at any NE, the expected amount of flow that goes through any edge of a min-cut set is equal to a constant fraction of its capacity. And from (23), if $\mathbf{P2}$'s equilibrium strategy only disrupts edges of one min-cut set, then the probability with which an edge is disrupted is constant for all the edges of that min-cut set, irrespective of the capacities of these edges. We can deduce the following corollary that directly follows from Prop. 5:

COROLLARY 2.

$$\forall (\sigma^{1*}, \sigma^{2*}) \in \Sigma, \forall \text{ min-cut set } E(\{\mathcal{S}, \mathcal{T}\}), \forall (i, j) \in E(\{\mathcal{S}, \mathcal{T}\}), \exists x \in \text{supp}(\sigma^{1*}) \text{ s.t. } x_{ij} > 0.$$

That is, for any NE and for any edge of a min-cut set, there exists a flow chosen with non-zero probability that passes through that edge.

We apply the previous results to the following example.

EXAMPLE 5. Consider the network in Fig. 4. The only min-cost max-flow is the flow x^* that sends 1 unit of flow through $\{s, 1, 3, t\}$, $\{s, 2, 3, t\}$ and $\{s, 2, 4, t\}$, and the only min-cut set is $\{(1, 2), (2, 3), (2, 4)\}$; see Fig. 7.




FIGURE 7. Min-cost max-flow (bold blue) and min-cut set attack (dotted red). The labels in the boxes represent the edges that are saturated by the min-cost max-flow.

In this example, the $s - t$ paths that induce the smallest transportation cost are the ones taken by x^* ; thus (A1) is satisfied. Lemma 3 tells us that the flows sent with positive probability in equilibrium only take paths taken by x^* . By combining this fact with Prop. 5, we conclude that in any equilibrium, the expected amount of flow in each of the paths $\{s, 1, 3, t\}$, $\{s, 2, 3, t\}$ and $\{s, 2, 4, t\}$ is $\frac{1}{p_2}$.

Besides, the only edges that are saturated by x^* are edges (1, 3), (2, 3) and (2, 4). From Prop. 4, we obtain that only these three edges can be disrupted with a nonzero probability in equilibrium. Hence, any $\mathbf{P2}$'s equilibrium strategy σ^{2*} is supported over at most $2^3 = 8$ pure actions instead of $2^9 = 512$ initial pure actions. These edges are exactly the min-cut set of the network in Fig. 7. Therefore, from Prop. 5, each of these edges is disrupted with probability $1 - \frac{3}{p_1}$.

In this example, we showed that the necessary conditions derived above help us restrict the pure actions that support equilibrium strategies to a significant extent. These properties also enable us to derive the following upper bounds on the probabilities with which actions in the support of $\tilde{\sigma}$ defined in Prop. 3 are chosen at any equilibrium.

PROPOSITION 6. Consider $(\sigma^{1*}, \sigma^{2*}) \in \Sigma$. Then we have the following bounds:

- (i) If $x^0 \in \text{supp}(\sigma^{1*})$, then $\sigma_{x^0}^{1*} \leq 1 - \frac{1}{p_2}$
- (ii) If $x^* \in \text{supp}(\sigma^{1*})$, then $\sigma_{x^*}^{1*} \leq \frac{1}{p_2}$
- (iii) If $\mu^{min} \in \text{supp}(\sigma^{2*})$, then $\sigma_{\mu^{min}}^{2*} \leq 1 - \frac{\alpha}{p_1}$
- (iv) If $\mu^0 \in \text{supp}(\sigma^{2*})$, then $\sigma_{\mu^0}^{2*} \leq \frac{\alpha}{p_1}$

Proof in Appendix A.3.

The NE $(\tilde{\sigma}^1, \tilde{\sigma}^2)$ derived in Prop. 3 attains these bounds. From these upper bounds, one can extend the game-theoretic intuition of Prop. 3 to any other NE. Specifically, when p_2 is close to 1, the probability with which x^0 can be chosen is very small. In contrast, when p_2 is large, x^* can be chosen only with small probability. Similarly, when p_1 is close to α , μ^{min} can be chosen only with small probability, and when p_1 is large, μ^0 can be chosen only with a small probability.

Lastly, we present a result analogous to the well-known Minimax Theorem by Von Neumann [20] for zero-sum games. Recall that the Minimax Theorem for a zero-sum game $\tilde{\Gamma} = \langle \{1, 2\}, (\mathcal{F}, \mathcal{A}), (\tilde{u}_1, -\tilde{u}_1) \rangle$ states that $\max_{\sigma^1 \in \Delta(\mathcal{F})} \min_{\sigma^2 \in \Delta(\mathcal{A})} \tilde{U}_1(\sigma^1, \sigma^2) = \min_{\sigma^2 \in \Delta(\mathcal{A})} \max_{\sigma^1 \in \Delta(\mathcal{F})} \tilde{U}_1(\sigma^1, \sigma^2)$ and $(\sigma^{1*}, \sigma^{2*}) \in \tilde{\Sigma} \iff \sigma^{1*} \in \arg \max_{\sigma^1 \in \Delta(\mathcal{F})} \min_{\sigma^2 \in \Delta(\mathcal{A})} \tilde{U}_1(\sigma^1, \sigma^2)$ and $\sigma^{2*} \in \arg \min_{\sigma^2 \in \Delta(\mathcal{A})} \max_{\sigma^1 \in \Delta(\mathcal{F})} \tilde{U}_1(\sigma^1, \sigma^2)$. Recall that the Minimax Theorem does not hold in general for games that are strategically equivalent to a zero-sum game. Interestingly, Γ satisfies some features of the Minimax Theorem.

PROPOSITION 7. Each player's payoffs for both maximinimizing and minimaximizing strategies are equal to the payoff at NE, i.e.,

$$\begin{aligned} \max_{\sigma^1} \min_{\sigma^2} U_1(\sigma^1, \sigma^2) &= 0 = \min_{\sigma^2} \max_{\sigma^1} U_1(\sigma^1, \sigma^2) \\ \max_{\sigma^2} \min_{\sigma^1} U_2(\sigma^1, \sigma^2) &= 0 = \min_{\sigma^1} \max_{\sigma^2} U_2(\sigma^1, \sigma^2) \end{aligned}$$

Furthermore, the set of minimaximizers is a superset of Σ , i.e., any NE is a minimaximizer.

Proof in Appendix A.3.

As in the Minimax Theorem, maximinimizing or minimaximizing each player's payoff gives the value of the game at equilibrium. In addition, Prop. 7 tells us that NE are minimaximizers. However, one can find minimaximizers that are not NE, which differs from the Minimax Theorem. Finally the proof of this proposition also implies that NE are not maximinimizers.

4. Budget-constrained game. From Thm. 1, we obtain that the *expected* cost of transportation (for **P1**) and the *expected* cost of attack (for **P2**) are constant in any NE. However, NE might differ from each other in the maximum cost of the actions chosen with positive probability. In this section, we view these costs as “budget expenditures” of the respective players. Recall the NE $(\tilde{\sigma}^1, \tilde{\sigma}^2)$ in Prop. 3, in which **P1** randomizes between x^0 and x^* , and in which **P2** randomizes between μ^0 and μ^{min} in Region III. To play the strategy $\tilde{\sigma}^1$ (resp. $\tilde{\sigma}^2$) in (11) (resp. (12)), **P1** (resp. **P2**) needs a budget of T^{min} (resp. F^{max}) for sending a min-cost max-flow (resp. attacking a min-cut set). We study the implications of the players not having a budget high enough to perform $(\tilde{\sigma}^1, \tilde{\sigma}^2)$, and more particularly we focus on computing the minimum budgets for which the results derived in Sections 3.2 and 3.3 still hold (with minor changes).

For **P1**, we use the infiniteness of her set of actions to find the lowest budget for which previous structural results hold. However, **P2**'s set of actions is discrete and we cannot derive the same bound. We approach this problem by restricting our attention to a subset of **P2**'s equilibrium strategies, specifically the strategies supported over the partitions of min-cut sets. Prop. 9 below

provides an explicit construction of such equilibrium strategies. Next, for this subset of equilibria, we formulate a problem for computing minimum budget equilibrium strategies as an integer programming problem. This provides a useful upper bound on **P2**'s lowest attack budget under which the previous results hold.

4.1. Revised model We now consider a revised game in which each player faces budget constraints noted b_1 and b_2 respectively: **P1** (resp. **P2**) can only send flows with transportation cost less than or equal to b_1 (resp. choose an attack with cost less than or equal to b_2). In this new game, the action sets \mathcal{F} and \mathcal{A} are revised to include the budget constraints b_1 and b_2 :

$$\mathcal{F}_{b_1} = \{x \in \mathcal{F} \mid T(x) \leq b_1\}, \quad \mathcal{A}_{b_2} = \{\mu \in \mathcal{A} \mid C(\mu) \leq b_2\}.$$

The new game is defined as $\Gamma_{b_1, b_2} = \langle \{1, 2\}, (\mathcal{F}_{b_1}, \mathcal{A}_{b_2}), (u_1, u_2) \rangle$ where u_1 (resp. u_2) is given in (2) (resp. (3)). As previously, we denote $\Delta(\mathcal{F}_{b_1})$ and $\Delta(\mathcal{A}_{b_2})$ the set of probability distributions over \mathcal{F}_{b_1} and \mathcal{A}_{b_2} . We also denote Σ_{b_1, b_2} the set of NE of the game Γ_{b_1, b_2} .

We investigate the question of computing the minimum budget b_1^* and b_2^* for **P1** and **P2** respectively for which the properties presented in Sections 3.2 and 3.3 hold for the NE of game $\Gamma_{b_1^*, b_2^*}$.

One way to tackle this problem is, given b_1 and b_2 , to find a NE of game Γ that satisfies the budget constraints. Note that $\mathcal{F}_{b_1} \subseteq \mathcal{F}$ and $\mathcal{A}_{b_2} \subseteq \mathcal{A}$. Therefore, if $\sigma^* \in \Sigma \cap (\Delta(\mathcal{F}_{b_1}) \times \Delta(\mathcal{A}_{b_2}))$, then $\sigma^* \in \Sigma_{b_1, b_2}$ (follows from (6) and (7)). In other words, a NE of game Γ that satisfies the budget constraints b_1 and b_2 is also a NE of game Γ_{b_1, b_2} . Consequently, with a little effort, one can show that the previous results are applicable to Γ_{b_1, b_2} .

Note that in the case when b_1 and b_2 are large enough so both players can send any flow in the network and attack any subset of edges (i.e., the budget constraints are not binding), $\mathcal{F}_{b_1} = \mathcal{F}$ and $\mathcal{A}_{b_2} = \mathcal{A}$, so $\Gamma_{b_1, b_2} = \Gamma$ and all the results derived in Section 3 are applicable.

We now focus on the more interesting case where the budget constraints are binding ($\mathcal{F}_{b_1} \subsetneq \mathcal{F}$ and $\mathcal{A}_{b_2} \subsetneq \mathcal{A}$). Thanks to the interchangeability of the NE, we can investigate each player's case independently while assuming that the other player's budget constraint is not binding.

4.2. P1's budget. In this subsection, we are looking to minimize the budget for transporting flows by **P1**, b_1^* , such that $\Sigma_{b_1^*, b_2}$ satisfies the properties presented in Sections 3.2 and 3.3. Without loss of generality, we assume that $b_2 \geq F^{\max} = C(\mu^{\min})$. This ensures that $\tilde{\sigma}^2$ in (12) is an equilibrium strategy for **P2** in $\Gamma_{b_1^*, b_2}$.

First, we argue that for Thm. 1 to hold for $\Gamma_{b_1^*, b_2}$, we need $b_1^* \geq \frac{1}{p_2} T^{\min}$. Indeed, if there existed a NE $(\sigma^{1*}, \sigma^{2*}) \in \Sigma_{b_1, b_2}$ with $b_1 < \frac{1}{p_2} T^{\min}$, then we would have:

$$\mathbb{E}_{\sigma^*} [T(x)] = \sum_{x \in \mathcal{F}_{b_1}} \sigma_x^{1*} T(x) < \frac{1}{p_2} T^{\min} \sum_{x \in \mathcal{F}_{b_1}} \sigma_x^{1*} = \frac{1}{p_2} T^{\min},$$

which contradicts (16). Therefore $b_1^* \geq \frac{1}{p_2} T^{\min}$.

Secondly, given any $b_1 \geq \frac{1}{p_2} T^{\min}$, we find an equilibrium strategy of Γ for **P1** that assigns positive probability on flows with transportation cost no greater than b_1 . This will ensure that $b_1^* \leq \frac{1}{p_2} T^{\min}$. For $b_1 \geq T^{\min}$, $x^* \in \mathcal{F}_{b_1}$ so $\tilde{\sigma}$ from Prop. 3 is a NE of Γ_{b_1, b_2} . However, for $\frac{1}{p_2} T^{\min} \leq b_1 \leq T^{\min}$, then $x^\dagger := \frac{b_1}{T^{\min}} x^* \in \mathcal{F}_{b_1}$ and the following proposition gives a NE of Γ_{b_1, b_2} .

PROPOSITION 8. *If $p_1 > \alpha$, $p_2 > 1$, $\frac{T^{\min}}{p_2} \leq b_1 \leq T^{\min}$, $b_2 \geq F^{\max}$, and under (A1), then $\exists \sigma^* = (\sigma^{1*}, \sigma^{2*}) \in \Sigma_{b_1, b_2}$ such that $U_1(\sigma^{1*}, \sigma^{2*}) = U_2(\sigma^{1*}, \sigma^{2*}) = 0$, and $\text{supp}(\sigma^{1*}) = \{x^0, x^\dagger\}$ and $\text{supp}(\sigma^{2*}) = \{\mu^0, \mu^{\min}\}$. The corresponding probabilities are given by:*

$$\sigma_{x^0}^{1*} = 1 - \frac{T^{\min}}{p_2 b_1}, \quad \sigma_{x^\dagger}^{1*} = \frac{T^{\min}}{p_2 b_1} \quad (24)$$

$$\sigma_{\mu^0}^{2*} = \frac{\alpha}{p_1}, \quad \sigma_{\mu^{min}}^{2*} = 1 - \frac{\alpha}{p_1}. \quad (25)$$

Proof in Appendix B.1.

The difference between Prop. 8 and Prop. 3 is that since x^* is too costly to send, **P1** sends only a fraction of x^* but with higher probability (so (15) in Thm. 1 is still satisfied). This is possible because **P1** has a continuous set of actions.

Thus, Prop. 8 ensures that the previous results still hold when **P1**'s budget is greater than or equal to $\frac{1}{p_2} T^{\min}$. This can be concluded by following the corresponding proofs with Prop. 8 as a starting point instead of Prop. 3. Therefore $b_1^* \leq \frac{1}{p_2} T^{\min}$, and we can conclude that $b_1^* = \frac{1}{p_2} T^{\min}$.

4.3. P2's budget. Now, we are looking for the lowest budget for **P2**, b_2^* , such that Σ_{b_1, b_2^*} satisfies the properties presented in Sections 3.2 and 3.3. Analogously to Section 4.2, we investigate this case without loss of generality assuming that $b_1 \geq T^{\min} = T(x^*)$. This ensures that $\tilde{\sigma}^1$ in (11) is an equilibrium strategy for **P1** in Γ_{b_1, b_2^*} .

Note that $b_2^* \geq F^{\max} - \frac{1}{p_1} T^{\min}$. Indeed, if there existed a NE $(\sigma^{1*}, \sigma^{2*}) \in \Sigma_{b_1, b_2}$ with $b_2 < F^{\max} - \frac{1}{p_1} T^{\min}$, then we would have:

$$\mathbb{E}_{\sigma^*} [C(\mu)] = \sum_{\mu \in \mathcal{A}_{b_2}} \sigma_{\mu}^{2*} C(\mu) < \left(F^{\max} - \frac{1}{p_1} T^{\min} \right) \sum_{\mu \in \mathcal{A}_{b_2}} \sigma_{\mu}^{2*} = F^{\max} - \frac{1}{p_1} T^{\min},$$

which contradicts (17). Therefore $b_2^* \geq F^{\max} - \frac{1}{p_1} T^{\min} = (1 - \frac{\alpha}{p_1}) F^{\max}$.

Unfortunately, this bound is seldom tight due to the finiteness of **P2**'s set of actions. For example, consider the network given in Fig. 4. We can notice that attacking any edge incurs a cost of at least 1. Thus, if $b_2 = (1 - \frac{\alpha}{p_1}) F^{\max} < 1$, then $\mathcal{A}_{b_2} = \{\mu^0\}$ (**P2** cannot attack) and the desired structural properties of NE may no longer hold.

So far, we know that $b_2^* \leq F^{\max}$ thanks to Prop. 3 (for any budget $b_2 \geq F^{\max}$, $\tilde{\sigma}^2$ is an equilibrium strategy for **P2**). We focus on computing a better upper bound on b_2^* by first considering a large subset of equilibrium strategies supported over the partitions of the min-cut sets. Then we find, in this subset of equilibrium strategies, the ones that require the lowest attack budget.

4.3.1. Partition-based equilibrium strategies. Throughout this subsection, we consider a min-cut set $E(\{\mathcal{S}, \mathcal{T}\})$. Let $\{e_1, \dots, e_N\}$ denote the edges that constitute the min-cut set, where N is the number of edges in $E(\{\mathcal{S}, \mathcal{T}\})$. Recall that μ^{min} is the attack that disrupts all the edges of $E(\{\mathcal{S}, \mathcal{T}\})$.

A partition of $\{e_1, \dots, e_N\}$ of size n is a set $\{\mathcal{P}_1, \dots, \mathcal{P}_n\}$ such that:

$$\forall (i, j) \in \llbracket 1, n \rrbracket^2 \mid i \neq j: \mathcal{P}_i \cap \mathcal{P}_j = \emptyset, \text{ and } \bigcup_{k \in \llbracket 1, n \rrbracket} \mathcal{P}_k = \{e_1, \dots, e_N\}.$$

DEFINITION 1. We say that the set of attacks $\{\mu^1, \dots, \mu^n\}$ is a *partition* of μ^{min} if there exists a partition $\{\mathcal{P}_1, \dots, \mathcal{P}_n\}$ of the min-cut set $\{e_1, \dots, e_N\}$ of size n such that $\forall k \in \llbracket 1, n \rrbracket$, $\mu^k = \mathbb{1}_{\mathcal{P}_k}$ is the attack that disrupts the edges of \mathcal{P}_k , i.e.:

$$\forall k \in \llbracket 1, n \rrbracket, \mu_{ij}^k = \begin{cases} 1 & \text{if } (i, j) \in \mathcal{P}_k, \\ 0 & \text{otherwise.} \end{cases}$$

Note that $\sum_{k=1}^n \mu^k = \mu^{min}$. The following proposition computes NE based on the partitions of μ^{min} .

PROPOSITION 9. If $p_1 > \alpha$, $p_2 > 1$, and under (A1), then for any partition $\{\mu^1, \dots, \mu^n\}$ of μ^{min} of size n , there exists a NE with support and corresponding probabilities given according to the following two regions:

(i) Region **III.a**: if $\alpha < p_1 < \frac{n\alpha}{n-1}$ and $p_2 > 1$, then $\{x^0, x^*\}$ and $\{\mu^1, \dots, \mu^n\} \cup \{\mu^0\}$ are the support of a NE $(\sigma^{1*}, \sigma^{2*}) \in \Sigma$ defined by:

$$\begin{aligned} & - \sigma_{x^0}^{1*} = 1 - \frac{1}{p_2}, \quad \sigma_{x^*}^{1*} = \frac{1}{p_2} \\ & - \forall k \in \llbracket 1, n \rrbracket, \sigma_{\mu^k}^{2*} = 1 - \frac{\alpha}{p_1}, \quad \sigma_{\mu^0}^{2*} = 1 - n \left(1 - \frac{\alpha}{p_1} \right) \end{aligned}$$

(ii) Region **III.b**: if $p_1 > \frac{n\alpha}{n-1}$ and $p_2 > 1$, then $\{x^0, x^*\}$ and $\{\mu^1, \dots, \mu^n\} \cup \{\mu^{min}\}$ are the support of a NE $(\sigma^{1*}, \sigma^{2*}) \in \Sigma$ defined by:

$$\begin{aligned} & - \sigma_{x^0}^{1*} = 1 - \frac{1}{p_2}, \quad \sigma_{x^*}^{1*} = \frac{1}{p_2} \\ & - \forall k \in \llbracket 1, n \rrbracket, \sigma_{\mu^k}^{2*} = \frac{\alpha}{p_1(n-1)}, \quad \sigma_{\mu^{min}}^{2*} = 1 - \frac{n\alpha}{p_1(n-1)} \end{aligned}$$

Proof in Appendix B.2.

These NE are illustrated in Figs. 8 and 9.





FIGURE 8. Support of partition-based equilibrium strategies in Regions **I-III**.

Given $p_1 > \alpha$ and $n \in \llbracket 1, N \rrbracket$, let us note $\Sigma_{p_1}^n$ the set of **P2**'s equilibrium strategies described by Prop. 9 whose support has a size equal to $n+1$ (the support is based on a partition of μ^{min} of size n and, depending on p_1 , whether includes μ^0 or μ^{min} according to Prop. 9). We denote the set of **P2**'s equilibrium strategies described by Prop. 9 (for a fixed p_1) as:

$$\Sigma_{p_1} := \bigcup_{n \in \llbracket 1, N \rrbracket} \Sigma_{p_1}^n. \quad (26)$$

Prop. 9 enables us to have an analytical expression of a large number of **P2**'s equilibrium strategies. Indeed, given any partition of a min-cut set, we can find a corresponding partition-based equilibrium strategy for **P2** thanks to Prop. 9. Since there are $2^N - 1$ such partitions, Σ_{p_1} contains $2^N - 1$ equilibrium strategies for **P2**.

In Fig. 8, we find again Regions **I** and **II** outlined in Props. 1 and 2. For Σ_{p_1} , Prop. 9 splits Region **III** into two subregions, where each subregion considers the partitions of μ^{min} in a specific


FIGURE 9. Probability distribution of **P2**'s partition-based equilibrium strategies.

manner. In Region **III.a** ($\frac{n\alpha}{n-1} > p_1$), an equilibrium strategy for **P2** randomizes over the partition $\{\mu^1, \dots, \mu^n\}$ and μ^0 . However, in Region **III.b** ($\frac{n\alpha}{n-1} < p_1$), an equilibrium strategy for **P2** randomizes over the same partition and μ^{min} . Intuitively, if **P2** partitions μ^{min} in too many components (i.e., $\frac{n\alpha}{n-1}$ decreases), then she assigns positive probability to the min-cut set attack. However, if she partitions μ^{min} in fewer components, then she chooses no attack action with a nonzero probability.

REMARK 2. The case $n = 1$ corresponds to attacking the whole min-cut set (it's a partition of size 1). When n tends to 1 from above, $\frac{n\alpha}{n-1} \rightarrow +\infty$. Therefore, if we draw Fig. 8 in the case $n = 1$, we find again Fig. 3 (Region **III.a** expands); thus, Prop. 3 is a particular case of Prop. 9.

EXAMPLE 6. Let us illustrate Prop. 9 with the example in Fig. 4. Recall that the only min-cut set is $\{(1, 3), (2, 3), (2, 4)\}$ and the only min-cost max-flow sends one unit of flow through each of the paths $\{s, 1, 3, t\}$, $\{s, 2, 3, t\}$ and $\{s, 2, 4, t\}$.

Let us consider one partition $\{\{(1, 3), (2, 3)\}, \{(2, 4)\}\}$ of the min-cut set. From this partition, we construct the corresponding attacks μ^1 that disrupts edges $(1, 3)$ and $(2, 3)$, and μ^2 that disrupts edge $(2, 4)$. Thus, $\{\mu^1, \mu^2\}$ is a partition of μ^{min} . The results obtained by applying Prop. 9 to this example are presented in Fig. 10.

FIGURE 10. NE described in Prop. 9 based on the partition $\{\mu^1, \mu^2\}$.

Now that we can analytically compute many more NE, we can try to find the equilibrium strategies among Σ_{p_1} that require the lowest budget.

4.3.2. Optimization problem. In the previous subsection, we saw that any partition of μ^{min} (along with μ^0 , μ^{min}) can be used to explicitly construct a subset of equilibrium strategies for **P2**. Specifically, an equilibrium based on such a partition can be mapped into one of the two regions (**III.a-b**) illustrated in Fig. 8.

Without loss of generality, let us consider a unique min-cut set $\{e_1, \dots, e_N\}$ consisting of N edges. With a slight abuse of notation, let us denote c_k the capacity of edge e_k (for all $k \in \llbracket 1, N \rrbracket$).

First, note the following:

$$\sigma^2 \in \Delta(\mathcal{A}_{b_2}) \iff \forall \mu \in \text{supp}(\sigma^2), C(\mu) \leq b_2 \iff \max_{\mu \in \text{supp}(\sigma^2)} C(\mu) \leq b_2.$$

That is, a strategy satisfies the budget constraint if and only if the maximum cost of conducting an attack chosen with positive probability is no greater than the resource budget.

Therefore, we want to find a strategy in Σ_{p_1} (defined in (26)) that minimizes the maximum cost of conducting an attack chosen with a positive probability, i.e.:

$$\arg \min_{\sigma^{2*} \in \Sigma_{p_1}} \max_{\mu \in \text{supp}(\sigma^{2*})} C(\mu) \quad (27)$$

The following proposition gives the answer.

PROPOSITION 10. *Among the NE listed in Prop. 9, a strategy that minimizes the budget needed is based on a partition of μ^{min} of size $n^* = \min \left\{ \left\lfloor \frac{p_1}{p_1 - \alpha} \right\rfloor, N \right\}$, and is obtained by solving the following integer-programming problem:*

$$\begin{aligned} (IP) \quad & \text{minimize} \quad z \\ & \text{subject to } z \geq \sum_{l=1}^N c_l y_{lk}, \forall k \in \llbracket 1, n^* \rrbracket \\ & \sum_{k=1}^{n^*} y_{lk} = 1, \forall l \in \llbracket 1, N \rrbracket \\ & y_{lk} \in \{0, 1\}, \forall (l, k) \in \llbracket 1, N \rrbracket \times \llbracket 1, n^* \rrbracket. \end{aligned}$$

Proof in Appendix B.2.

Given a partition $\{\mu^1, \dots, \mu^n\}$ of μ^{min} , the support of the corresponding strategy in Σ_{p_1} is this partition along with μ^{min} or μ^0 depending on p_1 . Let us note $\{\mathcal{P}_1, \dots, \mathcal{P}_n\}$ the corresponding partition of the min-cut set. For every $k \in \llbracket 1, n \rrbracket$, one may notice the following:

$$C(\mu^k) = \sum_{(i,j) \in \mathcal{P}_k} c_{ij} = \sum_{l=1}^N c_l \mu_{e_l}^k. \quad (28)$$

The cost of conducting attack μ^k is equal to the sum of the capacities of the edges of the min-cut set that μ^k disrupts. Therefore, the more **P2** partitions μ^{min} , the less number of edges of the min-cut set each μ^k disrupts. In Prop. 9, we saw that when $\frac{n\alpha}{n-1} > p_1$, **P2** randomizes over the partition and no attack, so the maximum attacking cost is induced by one of the elements of the partition. Thus, **P2** needs to increase n . However, when n increases, $\frac{n\alpha}{n-1}$ decreases and we saw that when $\frac{n\alpha}{n-1} < p_1$, then μ^{min} enters the support and the budget that is needed is F^{\max} (the capacity of the min-cut). Therefore, **P2** needs to increase n until $n^* = \max \left\{ n \in \llbracket 1, N \rrbracket \mid \frac{n\alpha}{n-1} \geq p_1 \right\} = \min \left\{ \left\lfloor \frac{p_1}{p_1 - \alpha} \right\rfloor, N \right\}$.

Knowing the optimal size of the partition of μ^{min} , we can find a partition of size n^* that minimizes the maximum attacking cost. Thanks to (28), one can see that this is equivalent to assigning N objects (the edges of the min-cut set) of value c_l each, into n^* bags such that the maximum value of the bags is minimized. This is the purpose of (IP).

The optimal value of (IP), z^* , gives a new upper bound on b_2^* (the results derived in Sections 3.2 and 3.3 hold when $b_2 \geq z^* \geq b_2^* \geq F^{\max} - \frac{1}{p_1} T^{\min}$). Depending on the min-cut set, z^* may be much smaller than F^{\max} , which was the previous upper bound on b_2^* obtained thanks to Prop. 3. In addition, the optimal solution of (IP), y_{lk}^* , gives us the corresponding way of partitioning μ^{\min} ($y_{lk}^* = 1$ if and only if edge e_l is disrupted by μ^k), and Prop. 9 derives the corresponding probabilities.

Let us illustrate Prop. 10 with an example.

EXAMPLE 7. Once again, consider the network given in Fig. 4 and assume that $p_1 = 5$. First, let us enumerate all the equilibrium strategies of Σ_{p_1} in Fig. 11:




FIGURE 11. Enumeration of the equilibrium strategies in Σ_{p_1} .

Fig. 11a shows that the NE of $\Sigma_{p_1}^1$ contains an attack that induces a cost of 3, while Figs. 11b, 11c, and 11d show that the NE of $\Sigma_{p_1}^2$ contain attacks that induce at most a cost of 2. However, Fig. 11e shows that the NE of $\Sigma_{p_1}^3$ contains an attack that also has a cost of 3. Therefore, **P2**'s equilibrium strategies of Σ_{p_1} that require the lowest budget are based on a partition of size 2, which corresponds to $\min \left\{ \left\lfloor \frac{5}{5-3} \right\rfloor, 3 \right\}$. For this example, the previous results hold when $b_2 \geq 2 = z^*$.

Thus, by combining Props. 9 and 10, we can compute a new upper bound on the lowest attack budget for **P2** to which we can apply the results derived in Sections 3.2 and 3.3.

5. Discussion. Let us discuss the implications of relaxing some assumptions of our model. First, we study a network that does not satisfy (A1) to see under which circumstances our results can still be applied. Next, we present a weaker assumption than (A1) for which the results derived in Section 3 still hold. We also discuss our assumption regarding the cost of attack: in the model studied in this article, we supposed that the cost of attacking an edge was proportional to its capacity. We now present an example network with a general cost of attack to study how some of the results derived in this article apply to such networks.

5.1. Relaxing Assumption 1. Consider the graph given in Fig. 12a and consider the game Γ with $p_1 = p_2 = 6$. The unique min-cost max-flow sends one unit of flow through each of the paths



FIGURE 12. Removal of paths that are too costly for **P1**.

$\{s, 1, t\}$ and $\{s, 2, t\}$ whose marginal transportation cost is equal to 5. However, $\{s, 1, 2, t\}$ has a marginal transportation cost equal to 3 so (A1) does not hold. We can show that $\tilde{\sigma}$ defined in (11) and (12) is not a NE anymore. Let us note x' the flow that sends 1 unit through path $\{s, 1, 2, t\}$, $\mu^1 = \mathbb{1}_{(s,1)}$, $\mu^2 = \mathbb{1}_{(1,2)}$ and $\mu^3 = \mathbb{1}_{(2,t)}$. Then, one can show that there exists an equilibrium where **P1**'s strategy σ^{1*} is defined by $\sigma_{x'}^{1*} = \frac{1}{6}$ and $\sigma_{x_0}^{1*} = \frac{5}{6}$, and **P2**'s strategy σ^{2*} is defined by $\sigma_{\mu^1}^{2*} = \sigma_{\mu^2}^{2*} = \sigma_{\mu^3}^{2*} = \frac{1}{6}$ and $\sigma_{\mu^0}^{2*} = \frac{1}{2}$. We can see that this strategy does not rely on the min-cost max-flow and the min-cut set anymore.

However, if we consider the game Γ for the parameter range $3 < p_1 < 5$ and $p_2 > 1$, then we can prove that $(\sigma^{1*}, \sigma^{2*})$ defined by $\sigma_{x_0}^{1*} = 1 - \frac{1}{p_2}$, $\sigma_{x'}^{1*} = \frac{1}{p_2}$, and $\sigma_{\mu^0}^{2*} = \frac{3}{p_1}$, $\sigma_{\mu^2}^{2*} = 1 - \frac{3}{p_1}$ is a NE. This result looks similar to the one we derived in Prop. 3. Actually, they are related: when $3 < p_1 < 5$, the marginal transportation costs of paths $\{s, 1, t\}$ and $\{s, 2, t\}$ are higher than the marginal value of effective flow, so **P1** has no incentive to send any flow along these paths. If we remove these paths from the graph (as in the elimination of strictly dominated strategies), we obtain the subgraph in Fig. 12b. It turns out this subgraph satisfies (A1) (it's only a path). Therefore, our results apply to this subgraph (i.e., the equilibrium we found is exactly $\tilde{\sigma}$ from Prop. 3 applied to this subgraph), and also to the original graph.

This example indicates that our results can be extended to networks for which (A1) holds for the reduced graph obtained after the removal of paths with transportation cost that are too large.

5.2. Transportation cost. It turns out the characterization of the NE of the game Γ (ref. Section 3) is also valid for a larger class of networks satisfying the following weaker assumption:

ASSUMPTION 2. *There exists an optimal solution of (\mathcal{P}_2) denoted $x^* \in \Omega^*$, and there exists a min-cut set $\{e_1, \dots, e_N\}$ with $\alpha_k := \min_{\{\lambda \in \Lambda \mid e_k \in \lambda\}} \sum_{(i,j) \in \lambda} b_{ij}$, $\forall k \in \llbracket 1, N \rrbracket$, such that for every $k \in \llbracket 1, N \rrbracket$, all $s-t$ paths taken by x^* that go through e_k have identical marginal transportation cost α_k , i.e.,*

$$\exists x^* \in \Omega^*, \exists \text{ min-cut set } \{e_1, \dots, e_N\} \mid \forall k \in \llbracket 1, N \rrbracket, \forall \lambda \in \Lambda \mid e_k \in \lambda: x_\lambda^* > 0 \implies \sum_{(i,j) \in \lambda} b_{ij} = \alpha_k.$$

In contrast to (A1), this assumption considers a class of networks whose min-cost max-flow takes paths that have different marginal transportation cost. Note that in (A2), when $\alpha_1 = \dots = \alpha_N \equiv \alpha$, we find again (A1). Below is an instance of a network that does not satisfy (A1) but satisfies (A2).

EXAMPLE 8. Consider the network flow problem in Fig. 13. There is a unique min-cost max-flow x^* , which carries 1 unit of flow through paths $\{s, 2, 4, t\}$, $\{s, 2, 3, t\}$ and $\{s, 1, t\}$. Thus, the total amount of flow is equal to 3 units. However, these paths induce different transportation costs; therefore (A1) is not satisfied.

Let us note $e_1 = (s, 1)$, $e_2 = (2, 3)$ and $e_3 = (4, t)$, then $\{e_1, e_2, e_3\}$ is a min-cut set. One can check that $\alpha_1 = 2$, $\alpha_2 = 3$ and $\alpha_3 = 4$. For $k \in \{1, 2, 3\}$, the paths taken by x^* that go through e_k induce a transportation cost equal to α_k . Thus, (A2) is satisfied, and with a little effort, we can check that our results in Section 3 also apply to the game Γ defined on this network.




FIGURE 13. Min-cost max-flow (bold blue) and min-cut set (dotted red) of a network satisfying (A2).

5.3. Cost of attack. Another assumption of our model is that the cost of attacking an edge is proportional to its capacity (ref. (3)). We investigate through an example the implications of considering the case of a more general cost structure for executing an attack. Let us consider the network given by Fig. 14, and once again note that $\tilde{\sigma}$ from (11) and (12) is not a NE anymore.




FIGURE 14. Network with general cost of attack. Edge capacities, transportation cost and cost of attack are denoted in red, green and orange colored labels respectively.

Let x^1 be the flow that sends 1 unit through path $\{s, 1, t\}$, x^2 be the flow that sends 1 unit through path $\{s, 2, t\}$, and $\mu' = \mathbb{1}_{\{(1,t), (2,t)\}}$ be the attack that disrupts edges $(1, t)$ and $(2, t)$. One can show that when $p_1 > 2$ and $p_2 > 3$, $(\sigma^{1*}, \sigma^{2*})$ defined by $\sigma_{x^0}^{1*} = 1 - \frac{3}{p_2}$, $\sigma_{x^1}^{1*} = \frac{1}{p_2}$, $\sigma_{x^2}^{1*} = \frac{2}{p_2}$ and $\sigma_{\mu^0}^{2*} = \frac{2}{p_1}$, $\sigma_{\mu'}^{2*} = 1 - \frac{2}{p_1}$ is a NE.

Notice that $\{(1, t), (2, t)\}$ is the cut-set of the graph that induces the smallest cost of attack. Hence, according to this NE, we conjecture that for networks with such general attack costs, the cut-set that induces the smallest attacking cost might form the support of an equilibrium strategy for **P2** (in contrast to the min-cut set in our article). Indeed, in our model, since the attacking cost

was proportional to the edge capacity, the min-cut sets were the cut-sets that induced the smallest cost of attack.

Regarding **P1**'s strategy, we can see that in this case, she routes x^2 twice as frequently as x^1 because the cost of attacking $(2, t)$ is twice as the cost of attacking $(1, t)$. This result differs from the NE we found in our model: each path taken by the min-cost max-flow was taken with the same probability, $\frac{1}{p_2}$, at equilibrium.

Appendix A: Proof of properties of Nash equilibria.

A.1. Proof of preliminary results.

PROOF OF LEMMA 1. u_1 is not an affine variant of $-u_2$. Indeed, let us suppose the contrary:

$$\exists (\lambda, \beta) \in \mathbb{R}_+^* \times \mathbb{R} \mid \forall (x, \mu) \in \mathcal{F} \times \mathcal{A}, u_1(x, \mu) = -\lambda u_2(x, \mu) + \beta$$

Noting (2) and (3), we have the following contradiction:

$$0 = u_1(x^0, \mu^0) + \lambda u_2(x^0, \mu^0) = \beta = u_1(x^*, \mu^0) + \lambda u_2(x^*, \mu^0) \neq 0.$$

Therefore, u_1 is not an affine variant of $-u_2$ and Γ is not an SCG (and a fortiori not a zero-sum game either).

However, the following transformations preserve the set of NE:

$$\frac{1}{p_1} u_1(x, \mu) + \frac{1}{p_2} C(\mu) = F(x^\mu) - \frac{1}{p_1} T(x) + \frac{1}{p_2} C(\mu) = \tilde{u}_1 \quad (29)$$

$$\frac{1}{p_2} u_2(x, \mu) - F(x) + \frac{1}{p_1} T(x) = -F(x^\mu) + \frac{1}{p_1} T(x) - \frac{1}{p_2} C(\mu) = -\tilde{u}_1 \quad (30)$$

So Γ is strategically equivalent to a zero-sum game $\tilde{\Gamma}$ and $\Sigma = \tilde{\Sigma}$, where $\tilde{\Sigma}$ denotes the set of NE of $\tilde{\Gamma}$. □

PROOF OF LEMMA 2. Suppose that **P1** chooses a flow x containing a loop l , i.e., the flow $x_l > 0$ stays in the loop and never reaches t . Then, if we note $x - x_l$ the flow resulting from removing the part of x that goes through l , we obtain $u_1(x, \mu) < u_1(x - x_l, \mu)$ for any attack μ . Thus, any flow containing loops is strictly dominated. □

PROOF OF PROPOSITION 1. Let us characterize the strategies that survive the iterated elimination of strictly dominated strategies.

$$\forall (x, \mu) \in \mathcal{F} \times \mathcal{A}, u_1(x, \mu) = p_1 F(x^\mu) - T(x) \stackrel{(8)}{\leq} p_1 F(x^\mu) - \alpha F(x) \leq (p_1 - \alpha) F(x) \leq 0.$$

If $x \neq x^0$, then $\forall \mu \in \mathcal{A}, u_1(x, \mu) < 0 = u_1(x^0, \mu)$. Therefore $x \neq x^0$ is strictly dominated and cannot be in the support of any NE.

Since $\forall \mu \in \mathcal{A}, u_2(x^0, \mu) = p_2 F(x^0 - (x^0)^\mu) - C(\mu) = -C(\mu)$, then, $\forall \mu \in \mathcal{A} \setminus \{\mu^0\}, u_2(x^0, \mu) = -C(\mu) < 0 = u_2(x^0, \mu^0)$. Hence, $\mu \neq \mu^0$ is now strictly dominated and cannot be in the support of any NE. Thus, (x^0, μ^0) is the unique strategy that survives the iterated elimination of strictly dominated strategies. We can conclude that (x^0, μ^0) is the unique NE when $p_1 < \alpha$. □

We need the following lemma to prove Prop. 2 and some of the subsequent results:

LEMMA 4. *The loss induced by any attack μ when a flow x is routed in the network is no greater than the cost of the attack μ .*

$$\forall (x, \mu) \in \mathcal{F} \times \mathcal{A}, \quad F(x - x^\mu) \leq C(\mu). \quad (31)$$

Furthermore, the loss induced by any attack μ that only disrupts edges of a min-cut set when a min-cost max-flow is routed in the network is equal to the cost of the attack μ .

$$\forall x^* \in \Omega^*, \quad \forall \mu \in \mathcal{A} \mid \exists \text{ min-cut set } E(\{\mathcal{S}, \mathcal{T}\}) \mid \forall (i, j) \in \mathcal{E}, \quad \mu_{ij} = 1 \implies (i, j) \in E(\{\mathcal{S}, \mathcal{T}\}):$$

$$F(x^* - x^{*\mu}) = C(\mu). \quad (32)$$

Proof of Lemma 4 Consider a flow $x \in \mathcal{F}$ and an attack $\mu \in \mathcal{A}$. First, notice the following:

$$F(x - x^\mu) = F(x) - F(x^\mu) = \sum_{\lambda \in \Lambda} x_\lambda - \sum_{\{\lambda \in \Lambda \mid \forall (i, j) \in \lambda, \mu_{ij} = 0\}} x_\lambda = \sum_{\{\lambda \in \Lambda \mid \exists (i, j) \in \lambda \mid \mu_{ij} = 1\}} x_\lambda, \quad (33)$$

because $\Lambda = \{\lambda \in \Lambda \mid \forall (i, j) \in \lambda, \mu_{ij} = 0\} \cup \{\lambda \in \Lambda \mid \exists (i, j) \in \lambda \mid \mu_{ij} = 1\}$.

Second, we have the following:

$$\begin{aligned} C(\mu) &= \sum_{\{(i, j) \in \mathcal{E} \mid \mu_{ij} = 1\}} c_{ij} \geq \sum_{\{(i, j) \in \mathcal{E} \mid \mu_{ij} = 1\}} x_{ij} \stackrel{(1)}{=} \sum_{\{(i, j) \in \mathcal{E} \mid \mu_{ij} = 1\}} \sum_{\{\lambda \in \Lambda \mid (i, j) \in \lambda\}} x_\lambda \\ &= \sum_{\{(i, j) \in \mathcal{E} \mid \mu_{ij} = 1\}} \sum_{\lambda \in \Lambda} x_\lambda \mathbb{1}_{\{(i, j) \in \lambda\}} = \sum_{\lambda \in \Lambda} x_\lambda \sum_{\{(i, j) \in \mathcal{E} \mid \mu_{ij} = 1\}} \mathbb{1}_{\{(i, j) \in \lambda\}} \end{aligned}$$

Notice that $\forall \lambda \in \Lambda \mid \forall (i, j) \in \lambda, \mu_{ij} = 0$, we have:

$$\sum_{\{(i, j) \in \mathcal{E} \mid \mu_{ij} = 1\}} \mathbb{1}_{\{(i, j) \in \lambda\}} = \sum_{\{(i, j) \in \lambda \mid \mu_{ij} = 1\}} 1 = 0,$$

and $\forall \lambda \in \Lambda \mid \exists (i, j) \in \lambda \mid \mu_{ij} = 1$, we have:

$$\sum_{\{(i, j) \in \mathcal{E} \mid \mu_{ij} = 1\}} \mathbb{1}_{\{(i, j) \in \lambda\}} = \sum_{\{(i, j) \in \lambda \mid \mu_{ij} = 1\}} 1 \geq 1.$$

Therefore, we obtain:

$$\begin{aligned} C(\mu) &\geq \sum_{\{\lambda \in \Lambda \mid \forall (i, j) \in \lambda, \mu_{ij} = 0\}} x_\lambda \times 0 + \sum_{\{\lambda \in \Lambda \mid \exists (i, j) \in \lambda \mid \mu_{ij} = 1\}} x_\lambda \times 1 \\ &\stackrel{(33)}{=} F(x - x^\mu) \end{aligned}$$

Now, consider a min-cost max-flow $x^* \in \Omega^*$ and an attack μ that disrupts only some edges of a min-cut set. If we follow the proof given previously, we can see that the inequalities become equalities in this case.

Indeed, $\forall (i, j) \in \mathcal{E} \mid \mu_{ij} = 1, c_{ij} = x_{ij}^*$ because the (min-cost) max-flows saturate all the edges of every min-cut set.

Besides, every path taken by a (min-cost) max-flow goes through only one edge of a min-cut set. Therefore:

$$C(\mu) = \sum_{\{\lambda \in \Lambda \mid \exists (i, j) \in \lambda \mid \mu_{ij} = 1\}} x_\lambda^* \sum_{\{(i, j) \in \lambda \mid \mu_{ij} = 1\}} 1 = \sum_{\{\lambda \in \Lambda \mid \exists (i, j) \in \lambda \mid \mu_{ij} = 1\}} x_\lambda^* = F(x^* - (x^*)^\mu)$$

□

PROOF OF PROPOSITION 2. First, $u_1(x^*, \mu^0) = p_1 F((x^*)^{\mu^0}) - T(x^*) = (p_1 - \alpha) F^{\max}$.

Second, $\forall x \in \mathcal{F}$, $u_1(x, \mu^0) = p_1 F(x^{\mu^0}) - T(x) \stackrel{(8)}{\leq} (p_1 - \alpha) F(x) \leq (p_1 - \alpha) F^{\max}$. So x^* is a BR for **P1**.

Similarly, $u_2(x^*, \mu^0) = p_2 F(x^* - (x^*)^{\mu^0}) - C(\mu^0) = 0$ because $(x^*)^{\mu^0} = x^*$. Besides,

$$\forall \mu \in \mathcal{A}, u_2(x^*, \mu) = p_2 F(x^* - (x^*)^\mu) - C(\mu) \leq F(x^* - (x^*)^\mu) - C(\mu) \leq 0$$

where the first inequality follows from $p_2 \leq 1$ and the second inequality follows from (31). Thus, μ^0 is a BR for **P2**.

Therefore $\forall x^* \in \Omega^*$, (x^*, μ^0) is a NE. □

PROOF OF PROPOSITION 3.

– First, let us show that any pure strategy is not a NE.

Let us suppose that **P2** chooses an attack μ that disrupts m edges e_1, \dots, e_m of \mathcal{G} (m can be equal to 0). Now, let us suppose that **P1** chooses a flow x that crosses one of the attacked edges, for instance e_1 . If we note x_p the part of x that goes through e_1 , then x_p will be lost because of the attack so the value of effective flow obtained by routing x is the same as if $x - x_p$ had been routed. However, the transportation cost of x is strictly greater than the transportation cost of $x - x_p$. Therefore $u_1(x, \mu) < u_1(x - x_p, \mu)$.

Thus, a BR for **P1** does not take paths containing at least one attacked edge. Let us note $\mathcal{G}^\mu = (\mathcal{V}, \mathcal{E} \setminus \{e_1, \dots, e_m\})$. Then, **P1**'s BR is a feasible flow in \mathcal{G}^μ . Now, let us suppose that **P1** chooses only such flows, the utility becomes $u_1(x, \mu) = p_1 F(x) - T(x)$. There are two cases:

- Case 1: there is no path in \mathcal{G}^μ with marginal transportation cost less than p_1 . Then, **P1**'s BR is x^0 (no flow). However, if **P1** chooses x^0 , then it's easy to see that **P2**'s BR is μ^0 (no attack), which means that the initial μ is not a BR for **P2** in this case. Notice that if the initial μ is μ^0 , then Case 1 is not satisfied (there are paths in \mathcal{G} of marginal transportation cost less than p_1 , because of the definition of α).
- Case 2: there exists at least one path in \mathcal{G}^μ with marginal transportation cost less than p_1 . Then, **P1**'s BR is to send as much flow as she can along the paths with maginal transportation cost less than p_1 (if there are different such flows with the same value, then **P1**'s BR is the one with least transportation cost). However, if **P1** chooses this BR, then the initial attack μ does not induce any loss, and **P2** will have an incentive to disrupt some edges of \mathcal{G}^μ instead (for instance we can prove that at least one edge is saturated by **P1**'s BR so **P2** will gain utility by attacking that edge since $p_2 > 1$). Thus, **P2**'s BR is different from μ .

Therefore, every pure strategy is not a NE.

– Now let us prove that $\tilde{\sigma} = (\tilde{\sigma}^1, \tilde{\sigma}^2)$ defined in (11) and (12) is a NE.

$$\forall \sigma^1 \in \Delta(\mathcal{F}), U_1(\sigma^1, \tilde{\sigma}^2) \stackrel{(4)}{=} p_1 \frac{\alpha}{p_1} \mathbb{E}_\sigma [F(x)] - \mathbb{E}_\sigma [T(x)] \stackrel{(8)}{\leq} \alpha \mathbb{E}_\sigma [F(x)] - \alpha \mathbb{E}_\sigma [F(x)] = 0.$$

$$\text{Besides, note that } U_1(\tilde{\sigma}^1, \tilde{\sigma}^2) = \alpha \frac{1}{p_2} F(x^*) - \frac{1}{p_2} T(x^*) \stackrel{(9)}{=} \frac{\alpha}{p_2} F(x^*) - \frac{\alpha}{p_2} F(x^*) = 0.$$

Similarly, note that

$$\begin{aligned} \forall \sigma^2 \in \Delta(\mathcal{A}), U_2(\tilde{\sigma}^1, \sigma^2) &\stackrel{(5)}{=} F(x^*) - \mathbb{E}_\sigma [F((x^*)^\mu)] - \mathbb{E}_\sigma [C(\mu)] \\ &= \mathbb{E}_\sigma [F(x^* - (x^*)^\mu) - C(\mu)] \stackrel{(31)}{\leq} 0, \end{aligned}$$

Finally, thanks to the Max-Flow Min-Cut Theorem, we have

$$U_2(\tilde{\sigma}^1, \tilde{\sigma}^2) = F(x^*) - \frac{\alpha}{p_1} F(x^*) - \left(1 - \frac{\alpha}{p_1}\right) C(\mu^{\min}) = 0.$$

Thus, $(\tilde{\sigma}^1, \tilde{\sigma}^2)$ is a NE. □

A.2. Proof of main theorem. We present two proofs of Thm. 1. The first proof involves the combination of the Max-Flow Min-Cut Theorem and best response inequalities. The second proof combines the Max-Flow Min-Cut Theorem with linear programming duality of $\tilde{\Gamma}$ defined in Lemma 1 (recall that $\Sigma = \tilde{\Sigma}$).

To prove Thm. 1, we need the following lemma:

LEMMA 5.

$$\forall(\sigma^{1*}, \sigma^{2*}) \in \Sigma, \forall x^* \in \Omega^*, \mathbb{E}_{\sigma^*} [F((x^*)^\mu)] = F(x^*) - \mathbb{E}_{\sigma^*} [C(\mu)]. \quad (34)$$

Proof of Lemma 5. We can relate both players' expected payoffs in two different ways:

$$U_1(\sigma^1, \sigma^2) = p_1 \mathbb{E}_\sigma [F(x)] - \mathbb{E}_\sigma [T(x)] - \frac{p_1}{p_2} \mathbb{E}_\sigma [C(\mu)] - \frac{p_1}{p_2} U_2(\sigma^1, \sigma^2) \quad (35)$$

$$U_2(\sigma^1, \sigma^2) = -\mathbb{E}_\sigma [C(\mu)] + p_2 \mathbb{E}_\sigma [F(x)] - \frac{p_2}{p_1} \mathbb{E}_\sigma [T(x)] - \frac{p_2}{p_1} U_1(\sigma^1, \sigma^2) \quad (36)$$

Let $\sigma^* = (\sigma^{1*}, \sigma^{2*}) \in \Sigma$. Since $(\tilde{\sigma}^1, \tilde{\sigma}^2) \in \Sigma$ (Prop. 3), we have:

$$0 = U_2(\tilde{\sigma}^1, \tilde{\sigma}^2) \stackrel{(7)}{\geq} U_2(\tilde{\sigma}^1, \sigma^{2*}) \stackrel{(5)}{=} F(x^*) - \mathbb{E}_{\sigma^*} [F((x^*)^\mu)] - \mathbb{E}_{\sigma^*} [C(\mu)]$$

So we get the following inequality:

$$\mathbb{E}_{\sigma^*} [F((x^*)^\mu)] \geq F(x^*) - \mathbb{E}_{\sigma^*} [C(\mu)] \quad (37)$$

Now, since $(\sigma^{1*}, \sigma^{2*}) \in \Sigma$:

$$\begin{aligned} U_1(\sigma^{1*}, \sigma^{2*}) &\stackrel{(6)}{\geq} U_1(\tilde{\sigma}^1, \sigma^{2*}) \stackrel{(4)}{=} \frac{p_1}{p_2} \mathbb{E}_{\sigma^*} [F((x^*)^\mu)] - \frac{1}{p_2} T(x^*) \\ &\stackrel{(9)}{=} \frac{p_1}{p_2} \mathbb{E}_{\sigma^*} [F((x^*)^\mu)] - \frac{\alpha}{p_2} F(x^*) \end{aligned} \quad (38)$$

By combining (7) and (35), and using $\tilde{\sigma}^2$ (ref. (12)), we obtain the following sequence of inequalities:

$$\begin{aligned} U_1(\sigma^{1*}, \sigma^{2*}) &\stackrel{(35)}{=} p_1 \mathbb{E}_{\sigma^*} [F(x)] - \mathbb{E}_{\sigma^*} [T(x)] - \frac{p_1}{p_2} \mathbb{E}_{\sigma^*} [C(\mu)] - \frac{p_1}{p_2} U_2(\sigma^{1*}, \sigma^{2*}) \\ &\stackrel{(8)}{\leq} (p_1 - \alpha) \mathbb{E}_{\sigma^*} [F(x)] - \frac{p_1}{p_2} \mathbb{E}_{\sigma^*} [C(\mu)] - \frac{p_1}{p_2} U_2(\sigma^{1*}, \sigma^{2*}) \\ &\stackrel{(7)}{\leq} (p_1 - \alpha) \mathbb{E}_{\sigma^*} [F(x)] - \frac{p_1}{p_2} \mathbb{E}_{\sigma^*} [C(\mu)] - \frac{p_1}{p_2} U_2(\sigma^{1*}, \tilde{\sigma}^2) \\ &\stackrel{(5)}{=} (p_1 - \alpha) \mathbb{E}_{\sigma^*} [F(x)] - \frac{p_1}{p_2} \mathbb{E}_{\sigma^*} [C(\mu)] - p_1 \mathbb{E}_{\sigma^*} [F(x)] \\ &\quad + \alpha \mathbb{E}_{\sigma^*} [F(x)] + \left(1 - \frac{\alpha}{p_1}\right) \frac{p_1}{p_2} C(\mu^{\min}) \end{aligned}$$

Therefore:

$$U_1(\sigma^{1*}, \sigma^{2*}) \leq \frac{p_1}{p_2} (C(\mu^{min}) - \mathbb{E}_{\sigma^*} [C(\mu)]) - \frac{\alpha}{p_2} C(\mu^{min}) \quad (39)$$

Combining (38) and (39), and using the Max-Flow Min-Cut Theorem, we obtain the reverse inequality:

$$\mathbb{E}_{\sigma^*} [F((x^*)^\mu)] \leq F(x^*) - \mathbb{E}_{\sigma^*} [C(\mu)] \quad (40)$$

From (37) and (40) we conclude that:

$$\mathbb{E}_{\sigma^*} [F((x^*)^\mu)] = F(x^*) - \mathbb{E}_{\sigma^*} [C(\mu)]$$

□

Now we can prove Thm. 1.

FIRST PROOF OF THEOREM 1. Let $\sigma^* = (\sigma^{1*}, \sigma^{2*}) \in \Sigma$. The first step is to show (iii) (the expected cost of attack), but we will need a few intermediate equations before.

Let us prove that $U_1(\sigma^{1*}, \sigma^{2*}) \geq 0$:

$$U_1(\sigma^{1*}, \sigma^{2*}) \stackrel{(6)}{\geq} U_1(x^0, \sigma^{2*}) = 0 \quad (41)$$

By combining (39) and (41), we obtain:

$$\mathbb{E}_{\sigma^*} [C(\mu)] \leq \left(1 - \frac{\alpha}{p_1}\right) C(\mu^{min}) \quad (42)$$

In order to get the reverse inequality, let us consider the strategy σ_ϵ^1 defined by $\sigma_{x^*}^1 = \frac{1+\epsilon}{p_2}$ and $\sigma_{x^0}^1 = 1 - \frac{1+\epsilon}{p_2}$ for an ϵ small enough (we can find such an ϵ and still have a probability distribution since $p_2 > 1$):

$$\begin{aligned} U_1(\sigma^{1*}, \sigma^{2*}) &\stackrel{(6)}{\geq} U_1(\sigma_\epsilon^1, \sigma^{2*}) \stackrel{(4)}{=} \frac{p_1(1+\epsilon)}{p_2} \mathbb{E}_{\sigma^*} [F((x^*)^\mu)] - \frac{1+\epsilon}{p_2} T(x^*) \\ &\stackrel{(9)}{=} \frac{p_1(1+\epsilon)}{p_2} \mathbb{E}_{\sigma^*} [F((x^*)^\mu)] - \frac{\alpha(1+\epsilon)}{p_2} F(x^*) \end{aligned}$$

Applying (34) from Lemma 5 gives us:

$$U_1(\sigma^{1*}, \sigma^{2*}) \geq \frac{p_1(1+\epsilon)}{p_2} (F(x^*) - \mathbb{E}_{\sigma^*} [C(\mu)]) - \frac{\alpha(1+\epsilon)}{p_2} F(x^*) \quad (43)$$

Now we combine (39) and (43) to get the following inequality:

$$\frac{p_1\epsilon}{p_2} F(x^*) - \frac{p_1\epsilon}{p_2} \mathbb{E}_{\sigma^*} [C(\mu)] - \frac{\alpha\epsilon}{p_2} F(x^*) \leq 0$$

which is equivalent to the desired inequality:

$$\mathbb{E}_{\sigma^*} [C(\mu)] \geq \left(1 - \frac{\alpha}{p_1}\right) F(x^*) \quad (44)$$

From (42), (44), and the Max-Flow Min-Cut Theorem, we obtain the expression of the expected cost of attack in any NE $\sigma^* \in \Sigma$:

$$\mathbb{E}_{\sigma^*} [C(\mu)] = \left(1 - \frac{\alpha}{p_1}\right) C(\mu^{min}) = \left(1 - \frac{\alpha}{p_1}\right) F^{\max} \stackrel{(9)}{=} F^{\max} - \frac{1}{p_1} T^{\min} \quad (45)$$

We can now use this equation in order to prove that $\mathbf{P1}$'s payoff is equal to 0 at equilibrium: by combining (39) and (45), we obtain:

$$U_1(\sigma^{1*}, \sigma^{2*}) \leq 0 \quad (46)$$

Therefore, (41) and (46) give us that $\mathbf{P1}$'s equilibrium payoff is zero:

$$U_1(\sigma^{1*}, \sigma^{2*}) = 0 \quad (47)$$

Let us now prove (15) (the expected amount of initial flow). Analogously to the previous arguments, we first prove that $U_2(\sigma^{1*}, \sigma^{2*}) \geq 0$:

$$U_2(\sigma^{1*}, \sigma^{2*}) \stackrel{(7)}{\geq} U_2(\sigma^{1*}, \mu^0) = 0 \quad (48)$$

From our previous results:

$$\begin{aligned} U_2(\sigma^{1*}, \sigma^{2*}) &\stackrel{(36)}{=} -\mathbb{E}_{\sigma^*} [C(\mu)] + p_2 \mathbb{E}_{\sigma^*} [F(x)] - \frac{p_2}{p_1} \mathbb{E}_{\sigma^*} [T(x)] - \frac{p_2}{p_1} U_1(\sigma^{1*}, \sigma^{2*}) \\ &\stackrel{(47)}{=} -\mathbb{E}_{\sigma^*} [C(\mu)] + p_2 \mathbb{E}_{\sigma^*} [F(x)] - \frac{p_2}{p_1} \mathbb{E}_{\sigma^*} [T(x)] \\ &\stackrel{(8)}{\leq} -\mathbb{E}_{\sigma^*} [C(\mu)] + p_2 \left(1 - \frac{\alpha}{p_1}\right) \mathbb{E}_{\sigma^*} [F(x)] \\ &\stackrel{(45)}{=} \left(1 - \frac{\alpha}{p_1}\right) (p_2 \mathbb{E}_{\sigma^*} [F(x)] - C(\mu^{min})) \end{aligned} \quad (49)$$

By combining (48) and (49), we obtain:

$$\mathbb{E}_{\sigma^*} [F(x)] \geq \frac{1}{p_2} C(\mu^{min}) \quad (50)$$

To get the reverse inequality, consider the strategy σ_ϵ^2 defined by $\sigma_{\mu^0}^2 = \frac{\alpha - \epsilon}{p_1}$ and $\sigma_{\mu^{min}}^2 = 1 - \frac{\alpha - \epsilon}{p_1}$, for an ϵ small enough (we can find such an ϵ and still have a probability distribution since $p_1 > \alpha$):

$$\begin{aligned} U_1(\sigma^{1*}, \sigma^{2*}) &\stackrel{(35)}{=} p_1 \mathbb{E}_{\sigma^*} [F(x)] - \mathbb{E}_{\sigma^*} [T(x)] - \frac{p_1}{p_2} \mathbb{E}_{\sigma^*} [C(\mu)] - \frac{p_1}{p_2} U_2(\sigma^{1*}, \sigma^{2*}) \\ &\stackrel{(8)}{\leq} (p_1 - \alpha) \mathbb{E}_{\sigma^*} [F(x)] - \frac{p_1}{p_2} \mathbb{E}_{\sigma^*} [C(\mu)] - \frac{p_1}{p_2} U_2(\sigma^{1*}, \sigma^{2*}) \\ &\stackrel{(7)}{\leq} (p_1 - \alpha) \mathbb{E}_{\sigma^*} [F(x)] - \frac{p_1}{p_2} \mathbb{E}_{\sigma^*} [C(\mu)] - \frac{p_1}{p_2} U_2(\sigma^{1*}, \sigma_\epsilon^2) \\ &\stackrel{(5)}{=} (p_1 - \alpha) \mathbb{E}_{\sigma^*} [F(x)] - \frac{p_1}{p_2} \mathbb{E}_{\sigma^*} [C(\mu)] - p_1 \mathbb{E}_{\sigma^*} [F(x)] \\ &\quad + (\alpha - \epsilon) \mathbb{E}_{\sigma^*} [F(x)] + \left(1 - \frac{\alpha - \epsilon}{p_1}\right) \frac{p_1}{p_2} C(\mu^{min}) \\ &= \frac{p_1}{p_2} (C(\mu^{min}) - \mathbb{E}_{\sigma^*} [C(\mu)]) - \frac{\alpha - \epsilon}{p_2} C(\mu^{min}) - \epsilon \mathbb{E}_{\sigma^*} [F(x)] \\ &\stackrel{(45)}{=} \frac{p_1}{p_2} C(\mu^{min}) - \frac{p_1}{p_2} \left(1 - \frac{\alpha}{p_1}\right) C(\mu^{min}) - \frac{\alpha - \epsilon}{p_2} C(\mu^{min}) - \epsilon \mathbb{E}_{\sigma^*} [F(x)] \\ &= \frac{\epsilon}{p_2} C(\mu^{min}) - \epsilon \mathbb{E}_{\sigma^*} [F(x)], \end{aligned}$$

and applying (41) gives us:

$$0 \leq \frac{\epsilon}{p_2} C(\mu^{min}) - \epsilon \mathbb{E}_{\sigma^*} [F(x)]$$

Thus, we obtain the desired inequality:

$$\mathbb{E}_{\sigma^*} [F(x)] \leq \frac{1}{p_2} C(\mu^{min}) \quad (51)$$

From (50), (51), and the Max-Flow Min-Cut Theorem, we obtain the expression of the expected amount of initial flow:

$$\mathbb{E}_{\sigma^*} [F(x)] = \frac{1}{p_2} F(x^*) = \frac{1}{p_2} F^{\max} \quad (52)$$

Likewise, by combining (49) and (52), we get:

$$U_2(\sigma^{1*}, \sigma^{2*}) \leq 0 \quad (53)$$

Equations (48) and (53) give us that $\mathbf{P2}$'s equilibrium payoff is also zero:

$$U_2(\sigma^{1*}, \sigma^{1*}) = 0 \quad (54)$$

Thus, we proved (i).

Now, by combining (5), (45), (52), and (54), we easily show (iv):

$$\mathbb{E}_{\sigma^*} [F(x^\mu)] = \frac{1}{p_2} F^{\max} - \frac{1}{p_2} \left(1 - \frac{\alpha}{p_1}\right) F^{\max} = \frac{\alpha}{p_1 p_2} F^{\max} \stackrel{(9)}{=} \frac{1}{p_1 p_2} T^{\min} \quad (55)$$

Finally, by combining (4), (47), and (55), we can finish proving (ii):

$$\mathbb{E}_{\sigma^*} [T(x)] = p_1 \frac{\alpha}{p_1 p_2} F^{\max} = \frac{\alpha}{p_2} F^{\max} \stackrel{(9)}{=} \frac{1}{p_2} T^{\min}.$$

□

SECOND PROOF OF THEOREM 1. We know that in a zero-sum game, each player's payoff is constant for any NE. Prop. 3 tells us that $\tilde{\sigma} \in \tilde{\Sigma}$. Therefore:

$$\begin{aligned} \forall(\sigma^{1*}, \sigma^{2*}) \in \tilde{\Sigma}, \quad \tilde{U}_1(\sigma^{1*}, \sigma^{2*}) &= \tilde{U}_1(\tilde{\sigma}^1, \tilde{\sigma}^2) \\ &\stackrel{(10)}{=} \frac{1}{p_2} \frac{\alpha}{p_1} F(x^*) - \frac{1}{p_1} \frac{1}{p_2} T(x^{max}) + \frac{1}{p_2} \left(1 - \frac{\alpha}{p_1}\right) C(\mu^{min}) \\ &= \frac{1}{p_2} \left(1 - \frac{\alpha}{p_1}\right) F^{\max} \end{aligned} \quad (56)$$

and this quantity is the optimal value of (LP_1) .

Let $\sigma^* = (\sigma^{1*}, \sigma^{2*}) \in \Sigma$. First, let us prove again Lemma 5:

By interchangeability, since $\tilde{\sigma} \in \tilde{\Sigma}$, then:

$$\frac{1}{p_2} \left(1 - \frac{\alpha}{p_1}\right) F^{\max} \stackrel{(56)}{=} \tilde{U}_1(\tilde{\sigma}^1, \tilde{\sigma}^2) = \tilde{U}_1(\tilde{\sigma}^1, \sigma^{2*}) = \frac{1}{p_2} \mathbb{E}_{\sigma^*} [F((x^*)^\mu)] - \frac{\alpha}{p_1 p_2} F(x^*) + \frac{1}{p_2} \mathbb{E}_{\sigma^*} [C(\mu)]$$

which directly gives the result.

Now, we prove the equalities thanks to complementary slackness: recall that (LP_2) is the dual of (LP_1) (and vice-versa), which means that $\forall \mu \in \mathcal{A}$, σ_μ^2 is the dual variable associated with the

constraint $\tilde{U}_1(\sigma^1, \mu) \geq z$. Similarly, $\forall x \in \mathcal{F}$, σ_x^1 is the dual variable associated with the constraint $\tilde{U}_2(x, \sigma^2) \geq z'$. We know that the optimal value of (LP₁) (resp. (LP₂)) is $\frac{1}{p_2} \left(1 - \frac{\alpha}{p_1}\right) F^{\max}$ (resp. $-\frac{1}{p_2} \left(1 - \frac{\alpha}{p_1}\right) F^{\max}$). Therefore, since NE are the optimal solutions of these LPs, complementary slackness can be written as follows: $\forall(\sigma^{1*}, \sigma^{2*}) \in \tilde{\Sigma}$,

$$\forall x \in \mathcal{F}, \sigma_x^{1*} \left(\tilde{U}_2(x, \sigma^{2*}) + \frac{1}{p_2} \left(1 - \frac{\alpha}{p_1}\right) F^{\max} \right) = 0 \quad (57)$$

$$\forall \mu \in \mathcal{A}, \sigma_\mu^{2*} \left(\tilde{U}_1(\sigma^{1*}, \mu) - \frac{1}{p_2} \left(1 - \frac{\alpha}{p_1}\right) F^{\max} \right) = 0 \quad (58)$$

Prop. 3 tells us that $\mu^{\min} \in \text{supp}(\tilde{\sigma}^2)$ (or equivalently $\tilde{\sigma}_{\mu^{\min}}^2 > 0$), therefore, complementary slackness and interchangeability imply:

$$\begin{aligned} \frac{1}{p_2} \left(1 - \frac{\alpha}{p_1}\right) F^{\max} &\stackrel{(58)}{=} \tilde{U}_1(\sigma^{1*}, \mu^{\min}) \stackrel{(10)}{=} -\frac{1}{p_1} \mathbb{E}_{\sigma^*} [T(x)] + \frac{1}{p_2} F^{\max} \\ &\iff \mathbb{E}_{\sigma^*} [T(x)] = \frac{\alpha}{p_2} F^{\max} \stackrel{(9)}{=} \frac{1}{p_2} T^{\min} \end{aligned} \quad (59)$$

Likewise, $\mu^0 \in \text{supp}(\tilde{\sigma}^2)$ therefore we can complete proving (ii):

$$\begin{aligned} \frac{1}{p_2} \left(1 - \frac{\alpha}{p_1}\right) F^{\max} &\stackrel{(58)}{=} \tilde{U}_1(\sigma^{1*}, \mu^0) \stackrel{(10)}{=} \mathbb{E}_{\sigma^*} [F(x)] - \frac{1}{p_1} \mathbb{E}_{\sigma^*} [T(x)] \\ &\iff \mathbb{E}_{\sigma^*} [F(x)] \stackrel{(59)}{=} \frac{1}{p_2} \left(1 - \frac{\alpha}{p_1}\right) F^{\max} + \frac{\alpha}{p_1 p_2} F^{\max} = \frac{1}{p_2} F^{\max} \end{aligned} \quad (60)$$

Similarly, $x^* \in \text{supp}(\tilde{\sigma}^1)$ (or equivalently $\tilde{\sigma}_{x^*}^1 > 0$), therefore, complementary slackness and interchangeability imply:

$$\begin{aligned} -\frac{1}{p_2} \left(1 - \frac{\alpha}{p_1}\right) F^{\max} &\stackrel{(57)}{=} \tilde{U}_2(x^*, \sigma^{2*}) = -\tilde{U}_1(x^*, \sigma^{2*}) \stackrel{(10)}{=} -\mathbb{E}_{\sigma^*} [F((x^*)^\mu)] + \frac{1}{p_1} T(x^*) - \frac{1}{p_2} \mathbb{E}_{\sigma^*} [C(\mu)] \\ &\stackrel{(34)}{=} - (F^{\max} - \mathbb{E}_{\sigma^*} [C(\mu)]) + \frac{\alpha}{p_1} F^{\max} - \frac{1}{p_2} \mathbb{E}_{\sigma^*} [C(\mu)] \end{aligned}$$

which is equivalent to:

$$\left(1 - \frac{1}{p_2}\right) \mathbb{E}_{\sigma^*} [C(\mu)] = \left(1 - \frac{\alpha}{p_1}\right) F^{\max} - \frac{1}{p_2} \left(1 - \frac{\alpha}{p_1}\right) F^{\max} = \left(1 - \frac{1}{p_2}\right) \left(1 - \frac{\alpha}{p_1}\right) F^{\max}$$

Therefore, we obtain (iii) again:

$$\mathbb{E}_{\sigma^*} [C(\mu)] = \left(1 - \frac{\alpha}{p_1}\right) F^{\max} \stackrel{(9)}{=} F^{\max} - \frac{1}{p_1} T^{\min} \quad (61)$$

We can now deduce (iv):

$$\begin{aligned} \mathbb{E}_{\sigma^*} [F(x^\mu)] &\stackrel{(9)}{=} \tilde{U}_1(\sigma^{1*}, \sigma^{2*}) + \frac{1}{p_1} \mathbb{E}_{\sigma^*} [T(x)] - \frac{1}{p_2} \mathbb{E}_{\sigma^*} [C(\mu)] \\ &\stackrel{(56), (59), (61)}{=} \frac{1}{p_2} \left(1 - \frac{\alpha}{p_1}\right) F^{\max} + \frac{1}{p_1 p_2} F^{\max} - \frac{1}{p_2} \left(1 - \frac{\alpha}{p_1}\right) F^{\max} \\ &= \frac{\alpha}{p_1 p_2} F^{\max} \stackrel{(9)}{=} \frac{1}{p_1 p_2} T^{\min} \end{aligned}$$

Lastly, to show (i), we proceed as follows:

$$\begin{aligned}
U_1(\sigma^{1*}, \sigma^{2*}) &\stackrel{(29)}{=} p_1 \tilde{U}_1(\sigma^{1*}, \sigma^{2*}) - \frac{p_1}{p_2} \mathbb{E}_{\sigma^*} [C(\mu)] \stackrel{(56), (61)}{=} \frac{p_1}{p_2} \left(1 - \frac{\alpha}{p_1}\right) F^{\max} - \frac{p_1}{p_2} \left(1 - \frac{\alpha}{p_1}\right) F^{\max} = 0 \\
U_2(\sigma^{1*}, \sigma^{2*}) &\stackrel{(30)}{=} p_2 \tilde{U}_2(\sigma^{1*}, \sigma^{2*}) + p_2 \mathbb{E}_{\sigma^*} [F(x)] - \frac{p_2}{p_1} \mathbb{E}_{\sigma^*} [T(x)] \\
&\stackrel{(56), (59), (60)}{=} - \left(1 - \frac{\alpha}{p_1}\right) F^{\max} + F^{\max} - \frac{\alpha}{p_1} F^{\max} = 0
\end{aligned}$$

□

A.3. Proof of necessary conditions.

PROOF OF LEMMA 3. Let us consider a NE $\sigma^* = (\sigma^{1*}, \sigma^{2*}) \in \Sigma$. Since $(\tilde{\sigma}^1, \tilde{\sigma}^2)$ is also a NE (Prop. 3), then, by interchangeability, $(\sigma^{1*}, \tilde{\sigma}^2)$ is a NE as well. So, thanks to Thm. 1, we have $\forall x \in \text{supp}(\sigma^{1*}), 0 = U_1(x, \tilde{\sigma}^2) = \alpha F(x) - T(x)$ where the first equality follows from (13). Therefore, $\forall x \in \text{supp}(\sigma^{1*}), T(x) = \alpha F(x)$. Since every path has a marginal transportation cost at least equal to α , then the last equality entails that any flow in the support of a NE takes paths that induce a marginal transportation cost equal to α .

□

PROOF OF PROPOSITION 4. Let us consider $(\sigma^{1*}, \sigma^{2*}) \in \Sigma$. We know that $(\tilde{\sigma}^1, \sigma^{2*}) \in \Sigma$ too (where $\tilde{\sigma}^1$ is defined in (11)). Then we can deduce that $\forall \mu \in \text{supp}(\sigma^{2*}), 0 = U_2(\tilde{\sigma}^1, \mu) = F(x^* - (x^*)^\mu) - C(\mu)$, where the first equality is a consequence of (14). Therefore, the Max-Flow Min-Cut Theorem gives us: $\forall \mu \in \text{supp}(\sigma^{2*}), C(\mu) = F(x^* - (x^*)^\mu) \leq F(x^*) = C(\mu^{\min})$.

Besides, since $\forall \mu \in \text{supp}(\sigma^{2*}), C(\mu) = F(x^* - (x^*)^\mu)$, then this means that the cost of conducting an attack that is in the support of a NE is equal to the loss it induces to any min-cost max-flow. Since the loss induced by an attack is never greater than the cost of the attack (see (31)), it means that each edge disrupted by an attack in the support of a NE is saturated by every min-cost max-flow (see the proof of Lemma 4).

□

PROOF OF PROPOSITION 5. Consider a min-cut set $E(\{\mathcal{S}, \mathcal{T}\})$. Given $p_1 > \alpha$ and $p_2 > 1$, one can find a NE $(\sigma^{1\ddagger}, \sigma^{2\ddagger})$ such that $\forall (i, j) \in E(\{\mathcal{S}, \mathcal{T}\}), \mathbb{1}_{\{(i, j)\}} \in \text{supp}(\sigma^{2\ddagger})$.

Consider a NE $\sigma^* = (\sigma^{1*}, \sigma^{2*}) \in \Sigma$, we know by interchangeability that $(\sigma^{1*}, \sigma^{2\ddagger}) \in \Sigma$. Therefore:

$$\begin{aligned}
\forall (i, j) \in E(\{\mathcal{S}, \mathcal{T}\}), U_2(\sigma^{1*}, \mathbb{1}_{\{(i, j)\}}) &\stackrel{(14)}{=} 0 \iff p_2 \mathbb{E}_{\sigma^*} [F(x - x^{\mathbb{1}_{\{(i, j)\}}})] - C(\mathbb{1}_{\{(i, j)\}}) = 0 \\
&\iff \mathbb{E}_{\sigma^*} [F(x - x^{\mathbb{1}_{\{(i, j)\}}})] = \frac{c_{ij}}{p_2} \\
&\iff \mathbb{E}_{\sigma^*} [x_{ij}] = \frac{c_{ij}}{p_2}
\end{aligned}$$

Now, suppose that σ^{2*} does not contain attacks that disrupt edges outside of $E(\{\mathcal{S}, \mathcal{T}\})$. First, notice that an edge e is disrupted if and only if it is attacked by at least one attack. Therefore:

$$\mathbb{P}_{\sigma^*} (\{e \text{ is disrupted}\}) = \sum_{\mu \in \mathcal{A}} \sigma_\mu^{2*} \mathbb{1}_{\{\mu_e=1\}} = \sum_{\{\mu \in \mathcal{A} \mid \mu_e=1\}} \sigma_\mu^{2*}.$$

Similarly, one can find a NE $(\sigma^{1'}, \sigma^{2'}) \in \Sigma$ such that $\forall e \in E(\{\mathcal{S}, \mathcal{T}\})$, there is a flow $x^e \in \text{supp}(\sigma^{1*})$ that crosses $E(\{\mathcal{S}, \mathcal{T}\})$ only at edge e and that takes paths of marginal transportation cost equal to α . We know by interchangeability that $(\sigma^{1'}, \sigma^{2*})$ is also a NE. Therefore:

$$\begin{aligned}
\forall e \in E(\{\mathcal{S}, \mathcal{T}\}), U_1(x^e, \sigma^{2*}) &\stackrel{(13)}{=} 0 \\
&\iff p_1 \mathbb{E}_{\sigma^*} [F((x^e)^\mu)] - \alpha F(x^e) = 0
\end{aligned}$$

$$\begin{aligned}
&\iff \sum_{\mu \in \mathcal{A}} \sigma_{\mu}^{2*} F((x^e)^{\mu}) = \frac{\alpha}{p_1} F(x^e) \\
&\iff \sum_{\{\mu \in \mathcal{A} \mid \mu_e = 1\}} \underbrace{\sigma_{\mu}^{2*} F((x^e)^{\mu})}_{=0} + \sum_{\{\mu \in \mathcal{A} \mid \mu_e = 0\}} \underbrace{\sigma_{\mu}^{2*} F((x^e)^{\mu})}_{=F(x^e)} = \frac{\alpha}{p_1} F(x^e) \\
&\iff F(x^e) \sum_{\{\mu \in \mathcal{A} \mid \mu_e = 0\}} \sigma_{\mu}^{2*} = \frac{\alpha}{p_1} F(x^e) \\
&\iff \frac{\alpha}{p_1} = \sum_{\{\mu \in \mathcal{A} \mid \mu_e = 0\}} \sigma_{\mu}^{2*} = 1 - \sum_{\{\mu \in \mathcal{A} \mid \mu_e = 1\}} \sigma_{\mu}^{2*} \\
&\iff 1 - \frac{\alpha}{p_1} = \sum_{\{\mu \in \mathcal{A} \mid \mu_e = 1\}} \sigma_{\mu}^{2*} = \mathbb{P}_{\sigma^*}(\{e \text{ is disrupted}\})
\end{aligned}$$

□

PROOF OF PROPOSITION 6. Consider a NE $\sigma^* = (\sigma^{1*}, \sigma^{2*}) \in \Sigma$. First, let us derive the bound for x^0 :

$$\frac{1}{p_2} F(x^*) \stackrel{(15)}{=} \mathbb{E}_{\sigma^*} [F(x)] = \sum_{x \in \mathcal{F} \setminus \{x^0\}} \sigma_x^{1*} F(x) \leq F(x^*) \sum_{x \in \mathcal{F} \setminus \{x^0\}} \sigma_x^{1*} = (1 - \sigma_{x^0}^{1*}) F(x^*)$$

Therefore, $\sigma_{x^0}^{1*} \leq 1 - \frac{1}{p_2}$.

Secondly, let us derive the bound for any min-cost max-flow x^* :

$$\frac{1}{p_2} F(x^*) \stackrel{(15)}{=} \mathbb{E}_{\sigma^*} [F(x)] = \sigma_{x^*}^{1*} F(x^*) + \sum_{x \in \mathcal{F} \setminus \{x^*\}} \sigma_x^{1*} F(x) \geq \sigma_{x^*}^{1*} F(x^*)$$

Therefore, $\sigma_{x^*}^{1*} \leq \frac{1}{p_2}$.

Thirdly, for any min-cut set attack μ^{min} , we have:

$$\left(1 - \frac{\alpha}{p_1}\right) C(\mu^{min}) \stackrel{(17)}{=} \mathbb{E}_{\sigma^*} [C(\mu)] = \sigma_{\mu^{min}}^{2*} C(\mu^{min}) + \underbrace{\sum_{\mu \in \mathcal{A} \setminus \{\mu^{min}\}} \sigma_{\mu}^{2*} C(\mu)}_{\geq 0}$$

Therefore, $\sigma_{\mu^{min}}^{2*} \leq 1 - \frac{\alpha}{p_1}$.

Finally, we can derive a similar bound for μ^0 :

$$\begin{aligned}
\left(1 - \frac{\alpha}{p_1}\right) C(\mu^{min}) &\stackrel{(17)}{=} \mathbb{E}_{\sigma^*} [C(\mu)] = \sum_{\mu \in \mathcal{A} \setminus \{\mu^0\}} \sigma_{\mu}^{2*} C(\mu) \\
&\stackrel{(21)}{\leq} C(\mu^{min}) \sum_{\mu \in \mathcal{A} \setminus \{\mu^0\}} \sigma_{\mu}^{2*} = (1 - \sigma_{\mu^0}^{2*}) C(\mu^{min})
\end{aligned}$$

Therefore, $\sigma_{\mu^0}^{2*} \leq \frac{\alpha}{p_1}$.

□

PROOF OF PROPOSITION 7.

– First, let us prove that $\max_{\sigma^1} \min_{\sigma^2} U_1(\sigma^1, \sigma^2) = 0$ by directly exhibiting a maximinimizer.

$$\forall (\sigma^1, \sigma^2) \in \Delta(\mathcal{F}) \times \Delta(\mathcal{A}), U_1(\sigma^1, \sigma^2) \stackrel{(4)}{=} \underbrace{p_1 \mathbb{E}_{\sigma} [F(x^{\mu})]}_{\geq 0} - \mathbb{E}_{\sigma} [T(x)] \geq -\mathbb{E}_{\sigma} [T(x)]$$

which is independent of σ^2 . Therefore: $\forall \sigma^1 \in \Delta(\mathcal{F})$, $\min_{\sigma^2} U_1(\sigma^1, \sigma^2) \geq -\mathbb{E}_\sigma [T(x)]$.

Besides $\forall \sigma^1 \in \Delta(\mathcal{F})$, $U_1(\sigma^1, \mu^{min}) = -\mathbb{E}_\sigma [T(x)]$. Therefore:

$$\forall \sigma^1 \in \Delta(\mathcal{F}), \min_{\sigma^2 \in \Delta(\mathcal{A})} U_1(\sigma^1, \sigma^2) = -\mathbb{E}_\sigma [T(x)] \leq 0.$$

This inequality tells us that $\max_{\sigma^1} \min_{\sigma^2} U_1(\sigma^1, \sigma^2) \leq 0$. Now it is easy to see that x^0 is a maximinimizer of U_1 :

$$\max_{\sigma^1 \in \Delta(\mathcal{F})} \min_{\sigma^2 \in \Delta(\mathcal{A})} U_1(\sigma^1, \sigma^2) = \min_{\sigma^2 \in \Delta(\mathcal{A})} U_1(x^0, \sigma^2) = 0 \quad (62)$$

- Now, let us prove that $\min_{\sigma^2} \max_{\sigma^1} U_1(\sigma^1, \sigma^2) \leq 0$ using the definition of a NE: let $(\sigma^{1*}, \sigma^{2*}) \in \Sigma$, then: $\forall \sigma^1 \in \Delta(\mathcal{F})$, $0 = U_1(\sigma^{1*}, \sigma^{2*}) \geq U_1(\sigma^1, \sigma^{2*})$ where the equality follows from (13) and the inequality follows from (6). Therefore, $\max_{\sigma^1} U_1(\sigma^1, \sigma^{2*}) = 0$. Then:

$$\min_{\sigma^2 \in \Delta(\mathcal{A})} \max_{\sigma^1 \in \Delta(\mathcal{F})} U_1(\sigma^1, \sigma^2) \leq \max_{\sigma^1 \in \Delta(\mathcal{F})} U_1(\sigma^1, \sigma^{2*}) = 0 \quad (63)$$

- We can get the reverse inequality thanks to the following inequality: $\max_{\sigma^1} \min_{\sigma^2} U_1(\sigma^1, \sigma^2) \leq \min_{\sigma^2} \max_{\sigma^1} U_1(\sigma^1, \sigma^2)$ (inequality that is true for any function of two variables):

$$0 \stackrel{(62)}{=} \max_{\sigma^1 \in \Delta(\mathcal{F})} \min_{\sigma^2 \in \Delta(\mathcal{A})} U_1(\sigma^1, \sigma^2) \leq \min_{\sigma^2 \in \Delta(\mathcal{A})} \max_{\sigma^1 \in \Delta(\mathcal{F})} U_1(\sigma^1, \sigma^2) \quad (64)$$

Therefore, (63) and (64) lead to:

$$\min_{\sigma^2 \in \Delta(\mathcal{A})} \max_{\sigma^1 \in \Delta(\mathcal{F})} U_1(\sigma^1, \sigma^2) = 0 \quad (65)$$

- The proof for U_2 is similar, but in this case μ^0 is the maximinimizer of U_2 .
- Let us prove that any NE of Γ is a minimaximizer: let $(\sigma^{1*}, \sigma^{2*}) \in \Sigma$, then:

$$0 \stackrel{(65)}{=} \min_{\sigma^2 \in \Delta(\mathcal{A})} \max_{\sigma^1 \in \Delta(\mathcal{F})} U_1(\sigma^1, \sigma^2) \leq \max_{\sigma^1 \in \Delta(\mathcal{F})} U_1(\sigma^1, \sigma^{2*}) \stackrel{(6)}{=} U_1(\sigma^{1*}, \sigma^{2*}) \stackrel{(13)}{=} 0$$

Therefore, $\min_{\sigma^2} \max_{\sigma^1} U_1(\sigma^1, \sigma^2) = \max_{\sigma^1} U_1(\sigma^1, \sigma^{2*})$, i.e., σ^{2*} is a *minimaximizer* of U_1 .

A similar argument tells us that σ^{1*} is a *minimaximizer* of U_2 .

□

Appendix B: Proof of budget-constrained game.

B.1. Proof of P1's budget.

PROOF OF PROPOSITION 8. Suppose that $p_1 > \alpha$, $p_2 > 1$, $\frac{T^{\min}}{p_2} \leq b_1 \leq T^{\min}$, $b_2 \geq F^{\max}$. Let us show that σ^* as defined in (24) and (25) is a NE.

$$\forall \sigma^1 \in \Delta(\mathcal{F}_{b_1}), U_1(\sigma^1, \sigma^{2*}) \stackrel{(4)}{=} p_1 \frac{\alpha}{p_1} \mathbb{E}_\sigma [F(x)] - \mathbb{E}_\sigma [T(x)] \stackrel{(8)}{\leq} \alpha \mathbb{E}_\sigma [F(x)] - \alpha \mathbb{E}_\sigma [F(x)] = 0$$

$$\text{Besides: } U_1(\sigma^{1*}, \sigma^{2*}) = \alpha \frac{T^{\min}}{p_2 b_1} F\left(\frac{b_1}{T^{\min}} x^*\right) - \frac{T^{\min}}{p_2 b_1} T\left(\frac{b_1}{T^{\min}} x^*\right) \stackrel{(9)}{=} \frac{\alpha}{p_2} F(x^*) - \frac{\alpha}{p_2} F(x^*) = 0$$

Similarly:

$$\begin{aligned} \forall \sigma^2 \in \Delta(\mathcal{A}_{b_2}), U_2(\sigma^{1*}, \sigma^2) &\stackrel{(5)}{=} \frac{T^{\min}}{b_1} F\left(\frac{b_1}{T^{\min}} x^*\right) - \frac{T^{\min}}{b_1} \mathbb{E}_\sigma \left[F\left(\left(\frac{b_1}{T^{\min}} x^*\right)^\mu\right) \right] - \mathbb{E}_\sigma [C(\mu)] \\ &= \mathbb{E}_\sigma [F(x^* - (x^*)^\mu) - C(\mu)] \stackrel{(31)}{\leq} 0 \end{aligned}$$

Besides: $U_2(\sigma^{1*}, \sigma^{2*}) = F(x^*) - \frac{\alpha}{p_1} F(x^*) - \left(1 - \frac{\alpha}{p_1}\right) C(\mu^{min}) = 0$ thanks to the Max-Flow Min-Cut Theorem. Thus, $(\sigma^{1*}, \sigma^{2*})$ is a NE. □

B.2. Proof of P2's budget.

PROOF OF PROPOSITION 9. Region III.a: if $\alpha < p_1 < \frac{n\alpha}{n-1}$ and $p_2 > 1$, then let us prove that $(\{x^0, x^*\}, \{\mu^0, \mu^1, \dots, \mu^n\})$ is the support of a NE $(\sigma^{1*}, \sigma^{2*})$ where:

$$\begin{aligned} - \sigma_{x^0}^{1*} &= 1 - \frac{1}{p_2}, \quad \sigma_{x^*}^{1*} = \frac{1}{p_2} \\ - \forall k \in \llbracket 1, n \rrbracket, \sigma_{\mu^k}^{2*} &= 1 - \frac{\alpha}{p_1}, \quad \sigma_{\mu^0}^{2*} = 1 - n \left(1 - \frac{\alpha}{p_1} \right) \end{aligned}$$

Let us first show (i) that σ^{1*} and σ^{2*} are probability distributions. Then let us prove (ii) $U_1(x^0, \sigma^{2*}) = U_1(x^*, \sigma^{2*}) = 0$ and (iii) $\forall x \in \mathcal{F}, U_1(x, \sigma^{2*}) \leq 0$. Similarly we prove (iv) $U_2(\sigma^{1*}, \mu^0) = U_2(\sigma^{1*}, \mu^1) = \dots = U_2(\sigma^{1*}, \mu^n) = 0$ and (v) $\forall \mu \in \mathcal{A}, U_2(\sigma^{1*}, \mu) \leq 0$

(i) First, let us prove that σ^{1*} and σ^{2*} are probability distributions:

σ^{1*} is clearly a probability distribution since $p_2 > 1$.

$$- p_1 > \alpha \text{ so } \sigma_{\mu^k}^{2*} = 1 - \frac{\alpha}{p_1} \geq 0$$

$$- \sigma_{\mu^0}^{2*} = 1 - n \left(1 - \frac{\alpha}{p_1} \right) = \frac{p_1 - np_1 + n\alpha}{p_1} = \frac{(n-1)(\frac{n\alpha}{n-1} - p_1)}{p_1} \geq 0 \text{ because } p_1 \leq \frac{n\alpha}{n-1}$$

$$- \sum_{\mu \in \mathcal{A}} \sigma_{\mu}^{2*} = 1$$

So σ^{2*} is a probability distribution.

(ii) Let us prove that $U_1(x^0, \sigma^{2*}) = 0$.

Since $\forall \mu \in \mathcal{A}, u_1(x^0, \mu) = 0$, then $U_1(x^0, \sigma^{2*}) = 0$.

Now, let us prove that $U_1(x^*, \sigma^{2*}) = 0$.

$$\begin{aligned} U_1(x^*, \sigma^{2*}) &= \sum_{k=1}^n \left(1 - \frac{\alpha}{p_1} \right) \left(p_1 F((x^*)^{\mu^k}) - T(x^*) \right) \\ &\quad + \left(1 - n \left(1 - \frac{\alpha}{p_1} \right) \right) \left(p_1 F((x^*)^{\mu^0}) - T(x^*) \right) \\ &\stackrel{(9)}{=} (p_1 - \alpha) \sum_{k=1}^n F((x^*)^{\mu^k}) + (p_1 - n(p_1 - \alpha)) F(x^*) - \alpha F(x^*) \end{aligned}$$

We can decompose x^* into $\{x^1, \dots, x^N\}$ where each x^l is the part of x^* that goes through e_l of the min-cut set $E(\{\mathcal{S}, \mathcal{T}\})$:

$$\forall l \in \llbracket 1, N \rrbracket, \forall (i, j) \in \mathcal{E}, x_{ij}^l = \sum_{\lambda \in \Lambda_{ij}^{e_l}} x_{\lambda}^*$$

where $\Lambda_{ij}^{e_l} = \{\lambda \in \Lambda \mid (i, j) \in \lambda \text{ and } e_l \in \lambda\}$ is the set of paths that go through (i, j) and e_l .

Therefore:

$$\sum_{k=1}^n F((x^*)^{\mu^k}) = \sum_{k=1}^n \sum_{l=1}^N F((x^l)^{\mu^k})$$

Moreover, only one of the μ^k affects x^l . Therefore we get:

$$\sum_{k=1}^n F((x^l)^{\mu^k}) = (n-1) F(x^l)$$

If we sum over l , we get:

$$\sum_{k=1}^n F\left((x^*)^{\mu^k}\right) = \sum_{l=1}^N \sum_{k=1}^n F\left((x^l)^{\mu^k}\right) = (n-1) \sum_{l=1}^N F(x^l) = (n-1) F(x^*)$$

which leads to:

$$U_1(x^*, \sigma^{2*}) = (p_1 - \alpha)(n-1) F(x^*) + (p_1 - n(p_1 - \alpha)) F(x^*) - \alpha F(x^*) = 0$$

(iii) Now, let us prove that $\forall x \in \mathcal{F}$, $U_1(x, \sigma^{2*}) \leq 0$.

$$\begin{aligned} \forall x \in \mathcal{F}, U_1(x, \sigma^{2*}) &= \sum_{k=1}^n \left(1 - \frac{\alpha}{p_1}\right) \left(p_1 F(x^{\mu^k}) - T(x)\right) \\ &\quad + \left(1 - n\left(1 - \frac{\alpha}{p_1}\right)\right) \left(p_1 F(x^{\mu^0}) - T(x)\right) \\ &= (p_1 - \alpha) \sum_{k=1}^n F(x^{\mu^k}) + (p_1 - n(p_1 - \alpha)) F(x) - T(x) \\ &\stackrel{(8)}{\leq} (p_1 - \alpha) \sum_{k=1}^n F(x^{\mu^k}) + (p_1 - n(p_1 - \alpha)) F(x) - \alpha F(x) \end{aligned}$$

Likewise: $\sum_{k=1}^n F(x^{\mu^k}) = (n-1) F(x)$. Therefore:

$$U_1(x, \sigma^{2*}) \leq ((p_1 - \alpha)(n-1) + (p_1 - n(p_1 - \alpha)) - \alpha) F(x) = 0$$

So σ^{1*} is a BR for **P1**.

(iv) Similarly, let us prove that $U_2(\sigma^{1*}, \mu^0) = 0$.

Since $\forall x \in \mathcal{F}$, $u_2(x, \mu^0) = p_2 F(x - x^{\mu^0}) - C(\mu^0) = 0$, then $U_2(\sigma^{1*}, \mu^0) = 0$

Now, let us prove that $\forall k \in \llbracket 1, n \rrbracket$, $U_2(\sigma^{1*}, \mu^k) = 0$.

$$\forall k \in \llbracket 1, n \rrbracket, U_2(\sigma^{1*}, \mu^k) = F(x^* - (x^*)^{\mu^k}) - C(\mu^k) \stackrel{(32)}{=} 0$$

(v) Lastly, let us prove that $\forall \mu \in \mathcal{A}$, $U_2(\sigma^{1*}, \mu) \leq 0$.

$$\forall \mu \in \mathcal{A}, U_2(\sigma^{1*}, \mu) = F(x^* - (x^*)^\mu) - C(\mu) \stackrel{(31)}{\leq} 0$$

Thus, σ^{2*} is a BR for **P2**. Therefore, $(\sigma^{1*}, \sigma^{2*}) \in \Sigma$.

Region III.b: if $p_1 > \frac{n\alpha}{n-1}$ and $p_2 > 1$, then let us prove that $(\{x^0, x^*\}, \{\mu^1, \dots, \mu^n, \mu^{\min}\})$ is the support of a NE $(\sigma^{1*}, \sigma^{2*})$ where:

$$\begin{aligned} - \sigma_{x^0}^{1*} &= 1 - \frac{1}{p_2}, \quad \sigma_{x^*}^{1*} = \frac{1}{p_2} \\ - \sigma_{\mu^k}^{2*} &= \frac{\alpha}{p_1(n-1)} \quad \forall k \in \llbracket 1, n \rrbracket, \quad \sigma_{\mu^{\min}}^{2*} = 1 - \frac{n\alpha}{p_1(n-1)} \end{aligned}$$

Let us first prove (i) that σ^{1*} and σ^{2*} are probability distributions. Then let us prove (ii) $U_1(x^0, \sigma^{2*}) = U_1(x^*, \sigma^{2*}) = 0$ and (iii) $\forall x \in \mathcal{F}$, $U_1(x, \sigma^{2*}) \leq 0$. Similarly we prove (iv) $U_2(\sigma^{1*}, \mu^1) = \dots = U_2(\sigma^{1*}, \mu^n) = U_2(\sigma^{1*}, \mu^{\min}) = 0$ and (v) $\forall \mu \in \mathcal{A}$, $U_2(\sigma^{1*}, \mu) \leq 0$.

- (i) First, let us prove that σ^{1*} and σ^{2*} are probability distributions:

σ^{1*} is clearly a probability distribution since $p_2 > 1$.

$$- \sigma_{\mu^k}^{2*} = \frac{\alpha}{p_1(n-1)} \geq 0$$

$$- \sigma_{\mu^{min}}^{2*} = 1 - \frac{n\alpha}{p_1(n-1)} = \frac{1}{p_1} \left(p_1 - \frac{n\alpha}{n-1} \right) \geq 0 \text{ because } p_1 \geq \frac{n\alpha}{n-1}$$

$$- \sum_{\mu \in \mathcal{A}} \sigma_{\mu}^{2*} = 1$$

So σ^{2*} is a probability distribution.

- (ii) Let us prove that $U_1(x^0, \sigma^{2*}) = 0$.

Since $\forall \mu \in \mathcal{A}$, $u_1(x^0, \mu) = 0$, then $U_1(x^0, \sigma^{2*}) = 0$.

Now let us prove that $U_1(x^*, \sigma^{2*}) = 0$.

$$\begin{aligned} U_1(x^*, \sigma^{2*}) &= \sum_{k=1}^n \frac{\alpha}{p_1(n-1)} \left(p_1 F((x^*)^{\mu^k}) - T(x^*) \right) \\ &\quad + \left(1 - \frac{\alpha n}{p_1(n-1)} \right) \left(\underbrace{p_1 F((x^*)^{\mu^{min}})}_{=0} - T(x^*) \right) \\ &\stackrel{(9)}{=} \frac{\alpha}{n-1} \sum_{k=1}^n F((x^*)^{\mu^k}) - \alpha F(x^*) \end{aligned}$$

As previously, one can show that $\sum_{k=1}^n F((x^*)^{\mu^k}) = (n-1) F(x^*)$, which leads to:

$$U_1(x^*, \sigma^{2*}) = \frac{\alpha}{n-1} (n-1) F(x^*) - \alpha F(x^*) = 0$$

- (iii) Now let us prove that $\forall x \in \mathcal{F}$, $U_1(x, \sigma^{2*}) \leq 0$.

$$\begin{aligned} U_1(x, \sigma^{2*}) &= \sum_{k=1}^n \frac{\alpha}{p_1(n-1)} \left(p_1 F(x^{\mu^k}) - T(x) \right) \\ &\quad + \left(1 - \frac{\alpha n}{p_1(n-1)} \right) \left(\underbrace{p_1 F(x^{\mu^{min}})}_{=0} - T(x) \right) \\ &= \frac{\alpha}{n-1} \sum_{k=1}^n F(x^{\mu^k}) - T(x) \end{aligned}$$

$$\text{Likewise: } \sum_{k=1}^n F(x^{\mu^k}) = (n-1) F(x).$$

Therefore:

$$U_1(x, \sigma^{2*}) \stackrel{(8)}{\leq} \frac{\alpha}{n-1} (n-1) F(x) - \alpha F(x) = 0$$

So σ^{1*} is a BR for **P1**.

- (iv) Similarly, let us prove that $U_2(\sigma^{1*}, \mu^{min}) = 0$.

$$\begin{aligned} U_2(\sigma^{1*}, \mu^{min}) &= \left(1 - \frac{1}{p_2} \right) (p_2 \times 0 - C(\mu^{min})) + \frac{1}{p_2} \left(p_2 F(x^* - (x^*)^{\mu^{min}}) - C(\mu^{min}) \right) \\ &= F(x^*) - C(\mu^{min}) = 0 \end{aligned}$$

where the last equality follows from the Max-Flow Min-Cut Theorem.

Now, let us prove that $\forall k \in \llbracket 1, n \rrbracket$, $U_2(\sigma^{1*}, \mu^k) = 0$

$$\begin{aligned} \forall k \in \llbracket 1, n \rrbracket, U_2(\sigma^{1*}, \mu^k) &= \left(1 - \frac{1}{p_2}\right) (p_2 \times 0 - C(\mu^k)) + \frac{1}{p_2} (p_2 F(x^* - (x^*)^{\mu^k}) - C(\mu^k)) \\ &= F(x^* - (x^*)^{\mu^k}) - C(\mu^k) \stackrel{(32)}{=} 0 \end{aligned}$$

- (v) We already proved in Region **III.a** that: $\forall \mu \in \mathcal{A}$, $U_2(\sigma^{1*}, \mu) \leq 0$. Thus σ^{2*} is a BR for **P2**.
Therefore, $(\sigma^{1*}, \sigma^{2*}) \in \Sigma$. □

PROOF OF PROPOSITION 10. We want to solve the optimization problem given in (27). Notice that this optimization problem can be written as follows:

$$\min_{\sigma^{2*} \in \Sigma_{p_1}} \max_{\mu \in \text{supp}(\sigma^{2*})} C(\mu) = \min_{n \in \llbracket 1, N \rrbracket} \min_{\sigma^{2*} \in \Sigma_{p_1}^n} \max_{\mu \in \text{supp}(\sigma^{2*})} C(\mu). \quad (66)$$

Let us first start by finding $n^* \in \arg \min_{n \in \llbracket 1, N \rrbracket} \min_{\sigma^{2*} \in \Sigma_{p_1}^n} \max_{\mu \in \text{supp}(\sigma^{2*})} C(\mu)$. There are two cases to consider:

- Case 1: for $\frac{n\alpha}{n-1} < p_1$, then $\forall \sigma^{2*} \in \Sigma_{p_1}^n$, $\text{supp}(\sigma^{2*}) = \{\mu^1, \dots, \mu^n\} \cup \{\mu^{min}\}$ so:

$$\forall \sigma^{2*} \in \Sigma_{p_1}^n, \max_{\mu \in \text{supp}(\sigma^{2*})} C(\mu) = C(\mu^{min}) = F^{\max}.$$

Therefore, if $\frac{n\alpha}{n-1} < p_1$, then $\min_{\sigma^{2*} \in \Sigma_{p_1}^n} \max_{\mu \in \text{supp}(\sigma^{2*})} C(\mu) = F^{\max}$.

- Case 2: for $\frac{n\alpha}{n-1} \geq p_1$, then $\forall \sigma^{2*} \in \Sigma_{p_1}^n$, $\text{supp}(\sigma^{2*}) = \{\mu^1, \dots, \mu^n\} \cup \{\mu^0\}$ so:

$$\forall \sigma^{2*} \in \Sigma_{p_1}^n, \max_{\mu \in \text{supp}(\sigma^{2*})} C(\mu) = \max_{k \in \llbracket 1, n \rrbracket} C(\mu^k).$$

For any partition $\{\mu^1, \dots, \mu^n\}$ of μ^{min} , we note $\{\mathcal{P}_1, \dots, \mathcal{P}_n\}$ the corresponding partition of the min-cut set $\{e_1, \dots, e_N\}$ of capacities c_1, \dots, c_N . Then the cost of each μ^k is equal to the sum of the capacities of the edges it disrupts, i.e.:

$$\forall k \in \llbracket 1, n \rrbracket, C(\mu^k) = \sum_{e_l \in \mathcal{P}_k} c_l$$

Therefore:

$$\forall \sigma^{2*} \in \Sigma_{p_1}^n, \max_{\mu \in \text{supp}(\sigma^{2*})} C(\mu) = \max_{k \in \llbracket 1, n \rrbracket} \sum_{e_l \in \mathcal{P}_k} c_l.$$

One can see that the problem is equivalent to finding a partition $\{\mathcal{P}_1, \dots, \mathcal{P}_n\}$ of the min-cut set such that the maximum sum of the capacities of the edges constituting each element of the partition is minimized. n still being fixed, we want to solve the following bilevel optimization problem:

$$\psi(n) := \min_{\{\mathcal{P}_1, \dots, \mathcal{P}_n\}} \max_{k \in \llbracket 1, n \rrbracket} \sum_{e_l \in \mathcal{P}_k} c_l$$

Now, let us argue that the optimal value of the previous bilevel problem, $\psi(n)$, does not increase when partitioning the min-cut set into more pieces, i.e., if $n' > n$, then, $\psi(n') \leq \psi(n)$.

Indeed, let us consider $n \leq N - 1$, and let us note $\{\mathcal{P}_1^*, \dots, \mathcal{P}_n^*\}$ an optimal partitioning of the min-cut set of size n : $\psi(n) = \max_{k \in \llbracket 1, n \rrbracket} \sum_{e_l \in \mathcal{P}_k^*} c_l$.

Since $n \leq N - 1$, then at least one of the \mathcal{P}_k^* contains at least two edges (Dirichlet's principle). Without loss of generality, let us assume that \mathcal{P}_n^* contains at least two edges. Let us denote e_{k_0} one of the edges of \mathcal{P}_n^* .

Now, let us consider $\{\mathcal{P}_1, \dots, \mathcal{P}_{n+1}\}$ a partition of the min-cut set of size $n + 1$ such that $\forall k \in \llbracket 1, n - 1 \rrbracket$, $\mathcal{P}_k = \mathcal{P}_k^*$, $\mathcal{P}_n = \mathcal{P}_n^* \setminus \{e_{k_0}\}$ and $\mathcal{P}_{n+1} = \{e_{k_0}\}$.

Notice that $c_{k_0} \leq \sum_{e_l \in \mathcal{P}_n^*} c_l$ and $\sum_{e_l \in \mathcal{P}_n} c_l \leq \sum_{e_l \in \mathcal{P}_n^*} c_l$.

Thus, we constructed a partition of the min-cut set of size $n + 1$, such that:

$$\psi(n+1) \leq \max_{k \in \llbracket 1, n+1 \rrbracket} \sum_{e_l \in \mathcal{P}_k} c_l \leq \max_{k \in \llbracket 1, n \rrbracket} \sum_{e_l \in \mathcal{P}_k^*} c_l = \psi(n)$$

Therefore, ψ is a non-increasing function, which means that in (66), we need to increase n as much as possible. However, let us not forget that we are in the case where $\frac{n\alpha}{n-1} \geq p_1$.

Therefore, the optimal partitioning size in this case is $n^* = \max \left\{ n \in \llbracket 1, N \rrbracket \mid \frac{n\alpha}{n-1} \geq p_1 \right\} = \min \left\{ \left\lfloor \frac{p_1}{p_1 - \alpha} \right\rfloor, N \right\}$.

Finally, notice that for any partition $\{\mathcal{P}_1, \dots, \mathcal{P}_n\}$, we have

$$\forall k \in \llbracket 1, n \rrbracket, \sum_{e_l \in \mathcal{P}_k} c_l \leq \sum_{l=1}^N c_l = F^{\max},$$

where the equality follows from the Max-Flow Min-Cut Theorem.

This implies that $\psi(n^*) \leq F^{\max}$, which is the optimal value of the original problem in Case 1. Therefore n^* is the optimal partitioning size.

Now that we know n^* , we can derive an integer-programming problem that provides us with an optimal partition of μ^{\min} .

$$\begin{aligned} & \text{minimize} && z \\ & \text{subject to} && z \geq \sum_{l=1}^N c_l y_{lk}, \forall k \in \llbracket 1, n^* \rrbracket \\ & && \sum_{k=1}^{n^*} y_{lk} = 1, \forall l \in \llbracket 1, N \rrbracket \\ & && y_{lk} \in \{0, 1\}, \forall (l, k) \in \llbracket 1, N \rrbracket \times \llbracket 1, n^* \rrbracket. \end{aligned}$$

One can see that this integer-programming problem gives the optimal value $\psi(n^*)$ and an optimal way of partitioning the min-cut set thanks to the y_{lk} . Indeed, for every edge e_l , and any set \mathcal{P}_k , $y_{lk} = 1$ means that edge e_l goes in set \mathcal{P}_k . Thanks to the constraint $\sum_{k=1}^{n^*} y_{lk} = 1$, $\forall l \in \llbracket 1, N \rrbracket$, each edge e_l goes in exactly one set \mathcal{P}_k , thus creating a partition of the min-cut set.

Then, from this partition of the min-cut set, one can partition μ^{\min} accordingly, and use Prop. 9 in order to derive the corresponding probabilities.

□

Acknowledgments This work was supported in part by FORCES (Foundations Of Resilient CybEr-Physical Systems), which receives support from the National Science Foundation (NSF award numbers CNS-1238959, CNS-1238962, CNS-1239054, CNS-1239166), NSF CAREER award CNS-1453126, and the AFRL LABLET - Science of Secure and Resilient Cyber-Physical Systems (Contract ID: FA8750-14-2-0180, SUB 2784-018400).

References

- [1] Acemoglu, Daron, Azarakhsh Malekian, Asuman Ozdaglar. 2013. Network security and contagion. Working Paper 19174, National Bureau of Economic Research. doi:10.3386/w19174. URL <http://www.nber.org/papers/w19174>.
- [2] Adler, Ilan, Constantinos Daskalakis, Christos H. Papadimitriou. 2009. A note on strictly competitive games. Stefano Leonardi, ed., *Internet and Network Economics, 5th International Workshop, WINE 2009, Rome, Italy, December 14-18, 2009. Proceedings, Lecture Notes in Computer Science*, vol. 5929. Springer, 471–474.
- [3] Alan Washburn, Kevin Wood. 1995. Two-person zero-sum games for network interdiction. *Operations Research* **43**(2) 243 – 251.
- [4] Avenhaus, Rudolf, Morton John Canty. 2009. Inspection games. Robert A. Meyers, ed., *Encyclopedia of Complexity and Systems Science*. Springer, 4855–4868. doi:10.1007/978-0-387-30440-3_287. URL http://dx.doi.org/10.1007/978-0-387-30440-3_287.
- [5] Avenhaus, Rudolf, Bernhard Von Stengel, Shmuel Zamir. 2002. Inspection games. R.J. Aumann, S. Hart, eds., *Handbook of Game Theory with Economic Applications*, vol. 3, chap. 51. Elsevier, 1947–1987. URL <https://ideas.repec.org/h/eee/gamchp/3-51.html>.
- [6] Ball, Michael O., Bruce L. Golden, Rakesh V. Vohra. 1989. Finding the most vital arcs in a network. *Oper. Res. Lett.* **8**(2) 73–76.
- [7] Baykal-Gürsoy, Melike, Zhe Duan, H. Vincent Poor, Andrey Garnaev. 2014. Infrastructure security games. *European Journal of Operational Research* **239**(2) 469–478.
- [8] Bertsimas, Dimitris, Ebrahim Nasrabadi, James B. Orlin. 2013. On the power of randomization in network interdiction. *CoRR* abs/1312.3478. URL <http://arxiv.org/abs/1312.3478>.
- [9] Bertsimas, Dimitris, Ebrahim Nasrabadi, Sebastian Stiller. 2013. Robust and adaptive network flows. *Operations Research* **61**(5) 1218–1242.
- [10] Cormican, Kelly J., David P. Morton, R. Kevin Wood. 1998. Stochastic Network Interdiction. *Operations Research* **46**(2).
- [11] Daskalakis, Constantinos, Paul W. Goldberg, Christos H. Papadimitriou. 2009. The complexity of computing a nash equilibrium. *SIAM Journal on Computing* **39**(1) 195–259. doi:10.1137/070699652. URL <http://dx.doi.org/10.1137/070699652>.
- [12] Edmonds, Jack, Richard M. Karp. 1972. Theoretical improvements in algorithmic efficiency for network flow problems. *J. ACM* **19**(2) 248–264. doi:10.1145/321694.321699. URL <http://doi.acm.org/10.1145/321694.321699>.
- [13] Ford, L. R., D. R. Fulkerson. 1956. Maximal flow through a network. *Canadian Journal of Mathematics* **8** 399–404.
- [14] Goldberg, Andrew V, Robert E Tarjan. 1989. Finding minimum-cost circulations by canceling negative cycles. *Journal of the ACM (JACM)* **36**(4) 873–886.
- [15] Goyal, Sanjeev, Adrien Vigier. 2014. Attack, defence, and contagion in networks. *Review of Economic Studies* **81**(4) 1518–1542.
- [16] Gueye, Assane, Jean C. Walrand, Venkat Anantharam. 2010. Design of network topology in an adversarial environment. Tansu Alpcan, Levente Buttyán, John S. Baras, eds., *Decision and Game Theory for Security - First International Conference, GameSec 2010, Berlin, Germany, November 22-23, 2010. Proceedings, Lecture Notes in Computer Science*, vol. 6442. Springer, 1–20.
- [17] Hong, Sunghoon, Myrna Wooders. 2010. Strategic Network Interdiction. Vanderbilt University Department of Economics Working Papers 1010, Vanderbilt University Department of Economics. URL <http://ideas.repec.org/p/van/wpaper/1010.html>.
- [18] Kalai, Ehud, Eitan Zemel. 1982. Totally balanced games and games of flow. *Mathematics of Operations Research* **7**(3) 476–478. doi:10.1287/moor.7.3.476. URL <http://dx.doi.org/10.1287/moor.7.3.476>.
- [19] Laszka, Aron, Assane Gueye. 2013. Quantifying network topology robustness under budget constraints: General model and computational complexity. Sajal K. Das, Cristina Nita-Rotaru, Murat Kantarcioglu,

-
- eds., *Decision and Game Theory for Security - 4th International Conference, GameSec 2013, Fort Worth, TX, USA, November 11-12, 2013. Proceedings, Lecture Notes in Computer Science*, vol. 8252. Springer, 154–174.
- [20] Neumann, John. 1928. Zur theorie der gesellschaftsspiele. *Mathematische Annalen* **100**(1) 295–320.
 - [21] Neumayer, S., G. Zussman, R. Cohen, E. Modiano. 2008. Assessing the impact of geographically correlated network failures. *Military Communications Conference, 2008. MILCOM 2008. IEEE*. 1–6. doi:10.1109/MILCOM.2008.4753111.
 - [22] Ratliff, H. Donald, G. Thomas Sicilia, S. H. Lubore. 1975. Finding the n most vital links in flow networks. *Management Science* **21**(5) 531–539. doi:10.1287/mnsc.21.5.531. URL <http://dx.doi.org/10.1287/mnsc.21.5.531>.
 - [23] Sullivan, Kelly M., J. Cole Smith. 2014. Exact algorithms for solving a euclidean maximum flow network interdiction problem. *Networks* **64**(2) 109–124. doi:10.1002/net.21561. URL <http://dx.doi.org/10.1002/net.21561>.
 - [24] Szeto, W.Y. 2013. Routing and scheduling hazardous material shipments: Nash game approach. *Transportmetrica B: Transport Dynamics* **1**(3) 237–260.
 - [25] Wollmer, R. 1964. Removing Arcs from a Network. *Operations Research* **12**(6) 934–940.
 - [26] Wood, R. Kevin. 1993. Deterministic network interdiction. *Mathematical and Computer Modelling* **17**(2) 1 – 18.